DOCUMENT RESUME

ED 462 374 SP 040 440

AUTHOR Schoenfeld, Alan H.
TITLE Toward a Theory of Teaching-in-Context.
PUB DATE 1998-07-07
NOTE 107p.; Some text may not reproduce well.
AVAILABLE FROM For full text: http://www.gse.berkeley.edu/Faculty/aschoenfeld/TechInContext/teaching-in-context.html.
PUB TYPE Opinion Papers (120) -- Reports - Descriptive (141)
EDRS PRICE MF01/PC05 Plus Postage.
DESCRIPTORS *Decision Making; Educational Theories; Elementary Secondary Education; Modeling (Psychology); Teacher Attitudes; Teacher Behavior; Teaching Methods; *Teaching Models
IDENTIFIERS Action Plans; Teacher Knowledge

ABSTRACT

This paper describes a university research program, the Teacher Model Group, which developed a theory of teaching in context, examining how and why teachers make specific decisions and take specific actions as they teach. The teacher model includes descriptions of teacher beliefs, goals, and knowledge. Two core assumptions underlying the model are that the activation levels of beliefs, goals, and knowledge at any moment will be assigned so that the highest priority beliefs, goals, and knowledge are consistent and mutually supportive and teacher actions are selected in a way consistent with the teacher's current highly activated beliefs, goals, and knowledge. The paper describes models and the modeling process, characterizes the general structure of the theory of teaching in context, and examines the means by which specific models of individual teachers engaged in the act of teaching are constructed. It describes representational forms that the Teacher Model Group uses for descriptions of teaching and presents four case studies involving a relatively new educator who taught traditionally, an experienced teacher who taught non-traditionally, and two very experienced teachers who modeled one lesson with ease and had trouble modeling another. The paper discusses prospects, strengths, and limitations of the model. (Contains 68 references.) (SM)

Reproductions supplied by EDRS are the best that can be made from the original document.

ED 462 374

TOWARD A THEORY OF TEACHING-IN-CONTEXT

ALAN H. SCHOENFELD

PERMISSION TO REPRODUCE AND
DISSEMINATE THIS MATERIAL HAS
BEEN GRANTED BY

A. H. SCHOENFELD

TO THE EDUCATIONAL RESOURCES
INFORMATION CENTER (ERIC)

1

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION
CENTER (ERIC)

☐ This document has been reproduced as received from the person or organization originating it

☐ Minor changes have been made to improve reproduction quality

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy

SP040440

BEST COPY AVAILABLE

TOWARD A THEORY OF TEACHING-IN-CONTEXT

Alan H. Schoenfeld
University of California, Berkeley

1. INTRODUCTION AND OVERVIEW

This paper describes an ambitious yet constrained research program. Since the early 1990s, the Teacher Model Group at Berkeley has been working on the development of a theory of teaching-in-context. Our intention is to provide a detailed theoretical account of how and why teachers do what they do "on line" - that is, while they are engaged in the act of teaching. This theoretical characterization of teaching is embodied in a model of the teaching process. The model describes, at a level of mechanism, the ways in which the teacher's goals, beliefs, and knowledge interact, resulting in the teacher's moment-to-moment decision-making and actions.

The purpose of this introduction is to contextualize the enterprise and to provide some motivation for it. What are we trying to explain, and what is beyond the scope of our efforts? What do we mean by a theory? By a model? How does the model that we offer differ from other models of teaching or decision-making? And, why should anybody care about work of this nature?

1.1 What the theory does (in brief).

Let me begin with the broad context and explain what we are trying to do. A central contribution of our work is its focus on explaining, at a very fine-grained level of detail, how and why teachers make specific decisions and take specific actions as they are engaged in teaching. There is, of course, a vast literature describing teachers' knowledge, behavior, and on-line decision-making. There are generally encompassing volumes such as the third edition of the *Handbook of Research on Teaching* (Wittrock, 1986), which includes a review chapter focusing on teachers' thought processes (Clark & Peterson, 1986). There is volume 20 of the *Review of Research in Education* (Darling-Hammond, 1994), which includes a collection of articles devoted to teachers' knowledge and practice. There is the *Handbook of Research on the Teaching and Learning of Mathematics* (Grouws, 1992), which includes reviews of classroom culture, the effects of teaching practices, teachers' beliefs, and teacher knowledge - see especially the reviews of teacher knowledge by Fennema and Franke (1992) and of teacher beliefs by Thompson (1992). There is the recent *Handbook of Educational Psychology* (Berliner & Calfee, 1996), which contains descriptions of teachers' beliefs and knowledge (Calderhead, 1996) and descriptive models of the teaching process (Borko &

Putnam, 1996). And there is more. In short, the ground we set out to plow has been well plowed by others before us. What can we hope to add?

The research cited immediately above identifies and elaborates the myriad factors that shape what teachers do in classrooms. We have learned from that literature (reviewed in Clark & Peterson, 1986; Calderhead, 1996) that teachers' knowledge, beliefs, and goals are critically important determinants of what teachers do and why they do it. We have learned the details of knowledge and knowledge organization (Shulman, 1986, 1987) and about the foundations of expertise (Berliner, 1994). We have learned about teacher beliefs (Thompson, 1992), and that (cf. Cohen, 1990) there is not a straightforward connection between beliefs and actions. We have learned about planning (Clark & Yinger, 1987; Yinger, 1977). These are the pieces of the puzzle, to be sure. But what we have not learned is how the pieces fit together. That is the question addressed here. I address it at a very fine level of detail, with a focus on mechanism - an explanation of precisely why teachers make particular choices at each point of instruction and precisely which beliefs, goals, and knowledge those decisions depend upon. In broad terms, this is question of cognitive modeling; the work described here fits within the tradition of cognitive studies of human thinking and problem solving. In more focused terms, this work leans very heavily on cognitive studies of teachers' planning and decision-making (see, e.g., Clark & Peterson, 1986; Clark & Yinger, 1987; Leinhardt & Greeno, 1986; Leinhardt, 1993; Shavelson, 1986; Yinger, 1977).

Consider what is taking place at any point in a lesson, from the teacher's perspective. Note that the teacher has carried into the classroom a substantial body of knowledge. This includes knowledge of the content, of the school environment, and of the students and his or her history with them. At a more fine-grained level, it also includes various routines, scripts, and schemata for dealing with classroom content and process. Similarly, the teacher has entered the classroom with a complex set of beliefs about the school, the students, and the content. The teacher has general goals and plans for the instruction and the students, and specific goals and plans for this lesson and its component parts. Moreover, many of these goals are linked to "action plans," mechanisms by which the teacher expects to achieve those goals. But now we are in the classroom, with the lesson in play. Something happens. It could be that a segment of the class has been concluded as planned. It could be that a student is in the midst of working a problem (either correctly or incorrectly) at the board. It could be that an unexpected issue has just arisen. The question, given our understanding of the teacher, the context, and the current constraints, is: *Something has happened. What will the teacher do next, and (more importantly) why?*

In abstract terms, the basic idea behind our response to this question is as follows. The model of the teacher includes descriptions of the teacher's beliefs, goals, and knowledge. At any given moment, each of these has a particular activation level, an indication of how important it is at that time. (N.B. For purposes of linguistic

variety, I shall use the terms "activation" and "priority" as synonyms.) Two core assumptions underlying the model are that (1) the activation levels of beliefs, goals, and knowledge at any moment will be assigned so that, if possible, the highest priority beliefs, goals, and knowledge are consistent and mutually supportive; and (2) actions taken by the teacher will be selected in a way to be consistent with the teacher's current highly activated beliefs, goals, and knowledge. That is (in no particular order): The teacher's high priority goals will be consistent with the teacher's high activation beliefs; the actions the teacher undertakes or decides to undertake will be consistent with those goals and beliefs; and the teacher's actions will draw upon related knowledge that has a high activation level. Now, when "something happens," the current balance can be affected. If a classroom event results in the strong activation of a particular belief, or goal, or some body of knowledge, then there will be a change in the activation level of beliefs, goals, and knowledge. This may in turn result in the choice of a new action plan.

To make the preceding description concrete, let me provide an example that foreshadows a detailed discussion given later in this paper. Part of the lesson that I describe in Section 6 consisted of a rather open-ended review of ways to select numbers to use in calculating a best value to represent a set of measurements. Then the teacher planned to discuss the three classical statistical measures of central tendency (mean, median, and mode). For both the review and new discussion, the teacher's plan was to use the following questioning strategy: Ask students if they can think of or remember a way to represent the given data set. Work with them to elaborate that idea, and when it has been elaborated successfully, ask if they can think of another method.

Things went according to plan for the first two rounds of the questioning strategy. A student mentioned "average," and the teacher pursued its definition. In response to the question "Is there another way . . ." another student mentioned that one number on their list of measurements showed up a large number of times. That observation led to a discussion of the mode. Once again, the teacher asked for suggestions. A student then said, "This is a little complicated but I mean it might work. . . ." and went on to propose, in somewhat ambiguous terms, a complex way to combine the numbers. This student comment is "the event," a possible cause for change of direction. The question: How will the teacher's beliefs, goals, and knowledge combine to shape the teacher's response, and what is that response likely to be?

A priori, there are any number of ways a teacher might respond to a comment of this nature. For example, one response might be an acknowledgment of the question but a deferral of the consideration of its substance. The teacher, who has a plan to follow, might try to follow up with the student after class or with the whole class at some later time. Or, perhaps, the teacher might address the student's comments directly, even if it means deviating from the lesson plan for a

substantial chunk of the lesson. At issue, then: can the model of this particular teacher, in this context, tell us what the teacher is likely to do?

In the specific case discussed in Section 6, we have extensive data about the teacher from his own writings, from interviews, and more. Our model of the teacher includes the fact that he believes that it is important to "honor student inquiry and initiative" if possible. In the model, the student's comment would result in that particular belief being given a very high activation level. So, if conditions are appropriate, the teacher might well establish the goal of pursuing the student's comment. By virtue of his subject matter knowledge, he knows that the idea proposed by the student is both germane and useful; it is worth bringing to the class's attention, even at the cost of some time. It is also early enough in the lesson that the teacher may still have a have a chance of covering all the intended material. In short, the conditions are appropriate, and he will decide to pursue the issue raised by the student. In terms of the model, the teacher's combination of beliefs and knowledge results in the creation of a new high priority goal: to explore the suggestion made by the student. Next, how will he pursue it? One of this teacher's frequently invoked pedagogical routines is a form of questioning called "reflective tosses," in which he takes issues at hand and poses them as questions for purposes of student elaboration. Hence the teacher is likely to ask the student to clarify the formula and/or ask the class what they think of the idea. (In fact, he does both.)

This example suggests what the theory of teaching-in-context is concerned with, and the kind of classroom behavior (on the part of the teacher) that we believe can be modeled. If one has a good understanding of the teacher's beliefs, goals, plans, and knowledge (which includes various pedagogical and content knowledge, the teacher's classroom routines, etc.) in a particular context, then one should be able to provide coherent, detailed explanations of what the teacher does, and why. The explanation of "why" will include a description of the (re-)prioritization of goals, beliefs, and knowledge that results in the teacher's choosing to do what he or she does, along a description of which intellectual resources (e.g., scripts or routines) the teacher will draw upon to implement that decision. The ultimate quality of the theory and the models that embody it will be judged by their explanatory power, their predictive power, and their scope (the range of cases that they fit well). For an extended discussion of these criteria see Section 2.

1.2 What the theory doesn't do.

I began this paper with the statement that our research program is ambitious but constrained. The ambition should now be clear. Let me, then, turn to the constraints. First, note that the theory focuses on the world of the classroom as seen by the teacher. While this represents a significant aspect of classroom reality, it is by no means all of it: there is the reality from the students' point of view, and there is what might be called the reality of the classroom as co-

constructed by teacher and students. The focus on the teacher may, ultimately, turn out to be a serious limitation, especially in classrooms where decision- making authority is broadly distributed. But, our experience in looking at many videotapes has been that, at least in classrooms where the teacher plays a large role (which is the case in most of the classrooms we have encountered), a focus on the classroom as seen through the teacher's lens explains a significant proportion of what takes place. I note that co-construction is not absent from the story that is told - it is represented as seen from the teacher's point of view, as the teacher works to make sense of what the students do and to react accordingly.

Second, I stress that this is not a theory of teaching in general, but a theory of *teaching-in-context* . I do not in any way wish to underestimate the importance of contextual factors - there is a large and fundamentally important literature that examines the factors that shape what takes place in classrooms (see, e.g., Grossman & Stodolsky, 1994, and Talbert, McLaughlin, & Rowan, 1993). What takes place in classrooms - indeed, what is even possible in classrooms - is shaped in fundamental ways by social, economic, and organizational, and curricular factors (to name just a few), and any complete theory of teaching must address these directly. But remember, our attempt is to explain how and why teachers do what they do *while engaged in the act of teaching.* Contextual factors, like the co-construction of the environment with students, are taken into account via the lens of the teacher's perception. That is, they manifest themselves in the way that the teacher views them. In modeling a teacher we ask what the teacher perceives the constraints and opportunities to be, and how those perceptions affects what he or she does. So, for example, the presence or absence of good curricular and support materials, administrative flexibility or rigidity, the perceived strengths and limitations of the students, as individuals and groups, are all in the model - as seen through the teacher's eyes. Our goal is the following. Given a sense of what the teacher sees (which includes perceptions of context and students), believes, knows, and wants to achieve, we want to be able to explain how those perceptions, beliefs, goals, and understandings combine to produce the teacher's decisions and actions as classroom interactions unfold.

1.3 Why do this kind of work?

The kind of work in which we are engaged is difficult and time-consuming. The theoretical perspective presented in this paper has taken years to evolve, and each of the case studies described, in which a specific teacher is modeled, has taken a substantial amount of time to develop. Simply put, why bother? What is the "value added" of this kind of work?

There are two kinds of reasons, theoretical and pragmatic. On the theoretical side, I will say simply that this kind of theory-building and modeling is a logical "next step" in developing an understanding of human cognition and problem solving. A quarter of a century ago, research was barely capable of capturing what happened in a twenty-minute problem-solving session in the laboratory -

indeed, the very dimensions of thinking and problem solving (the knowledge base, heuristic strategies, metacognition, beliefs) had yet to be established and elaborated. Hence laboratory studies, despite their "sterility" and remove from the chaos of classroom or real-world problem solving, were appropriate for the identification and elaboration of cognitive processes and for the development of a range of methodological tools. As our findings and methods have become more robust, we have been able to move from controlled settings into more complex and interactive ones. The goal, ultimately, is to develop a theory of human behavior in complex social settings. As noted above, there are strong constraints on what we model: we are trying to capture what the teacher perceives, thinks, and does in the "here and now," without taking into account broader contextual, sociological, or historical factors save as seen through the teacher's eyes. Nonetheless, building a fine-grained theory of how and why people do what they do in a social setting as complex as that of the classroom represents a significant advance toward more complete theoretical descriptions of human behavior. To our knowledge, the descriptions of teaching and their embodiments in detailed models at the level of mechanism given in this paper are the first of their kind.

The pragmatic reasons for doing this kind of work are straightforward. I shall discuss three. By way of preliminaries, let me stress that the model we have developed is deliberately not designed to focus on any one type of teaching. The structure of the model allows for accurate descriptions of a wide range of teaching - from the most traditional to the most "reform," from highly structured to free-form, and so on. The model focuses not only on why teachers do what they do (individual teachers' beliefs and goals), but on what enables them to do what they do - the various kinds of knowledge (subject matter knowledge, pedagogical knowledge, information stored in memory as schemata, scripts, routines, etc.) that undergird their classroom actions.

The first pragmatic justification for this kind of work is that the model can be used as an analytic tool to help understand and foster aspects of "good teaching." As stressed in the previous paragraph, the model itself does not take a value stance - it can be used to analyze any style of teaching, independent of whether the people doing the analysis happen to favor that style. But, one can of course choose to analyze teaching that one thinks is worth emulating. The model of any teacher describes the beliefs, goals, and knowledge that enable that teacher to do what he or she does. Hence, if you think that what a particular teacher does is noteworthy, you can use the model to examine the bases of that teacher's competence (exploring in detail which beliefs seem to be active, which decisions get made in which ways, which intellectual and material resources get accessed and used in which circumstances, and so on). This provides a sufficiency argument of sorts: if you understand the combination of things that enables that teacher to be successful at doing particular things, then there is a chance you can help people learn to achieve what that teacher does.

It should be stressed that the kind of analysis to which I refer is not at the level of behavior ("this is what the teacher does"). The analysis requires dealing with the whole package - the teacher's beliefs, goals, and knowledge. Yet, it is worth noting that apparent complexity can sometimes be deceiving and that a detailed analysis may reveal some surprisingly simple things. For example, we shall see later in this paper that Jim Minstrell and I, completely independently, developed parallel questioning styles much of whose essence can be "captured" with rather simple decision procedures. There is no suggestion that Minstrell or I consciously follow these decision procedures while we teach, or that we learned them directly at some point - indeed, there is no suggestion that either of us was necessarily aware of having such a strategy. Nonetheless, the decision procedures can be used as scaffolding to help other teachers learn to conduct question-based lessons in that manner. By way of analogy: I had tried, for many years and without success, to learn how to juggle. One day I chanced across a book (Lewis, 1974) in which the author had watched skilled jugglers, and "captured" what they do in a set of simple steps. There was no suggestion that the jugglers had learned to juggle that way, or that they were aware of the steps the author described. But, the description was accurate and productive. Soon after reading it I was able to juggle.

The model proposed here is complex. Some readers might be tempted to dismiss such things as being ivory tower pie-in-the-sky stuff, not relevant to classroom realities. Nothing could be further from the truth. Consider research on mathematical thinking and problem solving as a case in point. In essence, research on mathematical thinking from the 1970's through the 1990's identified the dimensions of mathematical performance that are necessary for successful problem solving. Prior to the 1970's the main focus of theory and practice had been on the knowledge base - facts, procedures, and conceptual understanding. Through the 1970's and 80's the scope of inquiry expanded to include problem solving strategies, metacognition, beliefs, and mathematical practices. It became understood that to be successful, a mathematics program must attend to such things - that students will develop problem solving skills, good metacognitive behaviors, and appropriate beliefs only when they learn mathematics in environments that support the development of such understandings. In fundamental ways, then, research into "what it takes to think mathematically" provided the theoretical underpinnings of the reform movement in mathematics instruction. Reform documents written for teachers did not generally emphasize the research base. But reform would not have been shaped the way it was, were it not for fundamental changes in our understanding of what constitute the underpinnings of mathematical thinking and problem solving. Likewise, an enhanced understanding of reform-based teaching may lead to different or expanded notions of what it takes to achieve various goals in the classroom.

The second justification for doing this kind of work is that the model provides a fine-grained analytical tool for examining the beliefs, goals, and knowledge of individual teachers. As such, it is a tool for professional development. A teacher

and others can explore what he or she was trying to achieve, what the options were, how well things did or didn't work - and in the process, reflect on ways that things might have been done differently. Of course, examining one's own teaching, say in "video clubs" (Frederiksen, Sipusic, Gamoran, & Wolfe, 1992), is an established form of professional development. The model offers an analytic tool with which the practice can be sharpened.

Third, the model can be used to trace the evolution (of beliefs, classroom goals, knowledge, etc.) of individual teachers. As such, it can be used to describe the "developmental trajectories" of individual teachers. Cumulatively, such data can be used to begin to paint portraits of what it means to evolve from a beginning teacher into an established professional. We know little about the evolution of professional competence, and such information would be useful for the field in general. What kinds of understandings tend to develop, and when? What kinds of targeted interventions might be most useful, at what points in people's careers? In short, the more refined an understanding we have of how teaching works, the more we will be in a position to promote the development of teaching competence. That said, let me preview the content of this paper.

Section 2 presents an extended discussion of models and the modeling process. The purpose of that discussion is to specify what I mean by a model and to describe some criteria by which the adequacy and usefulness of models can be judged. I think Section 2 is necessary, though not exciting reading - there is a good deal of confusion about the term "model," and I want to be very clear about the nature of the work in which we are engaged. Readers may want to move through it rapidly, to get to the substance of the paper.

Section 3 provides a characterization of the general structure of the theory of teaching-in-context. (The model, of course, is the concrete and testable embodiment of the theory.) There I delineate the components of the model and the mechanisms by which they interact. Then, in Section 4, I describe the means by which specific models of individual teachers engaged in the act of teaching are constructed. In order to lay the groundwork for the cases that follow, Section 4 describes the representational forms that the Teacher Model Group uses for its descriptions of teaching. With that as backdrop, I turn to four case studies. (I understand that for many readers these cases will be the main focus of the paper, and apologize for the delayed gratification. What precedes them should help illuminate the cases.)

The two main cases I discuss are given in Sections 5 and 6. There I present the fairly detailed analyses of two lesson segments that differ along a number of important dimensions. Section 5 describes a segment of a lesson taught by a relatively new teacher that covers rather traditional material from a standard textbook. Section 6 describes a decidedly non-traditional lesson taught by a very experienced teacher. It will be seen that the model accommodates these two very different lessons, taught by two very different instructors, without difficulty.

Equally important, the compare-and-contrast of what these teachers are able to do (and more importantly, what enables them to do it) raises interesting issues that are pursued in the concluding discussion. Section 7 offers brief descriptions of two additional cases - of two very experienced teachers who, like the teacher of the lesson described in Section 6, are known for their instruction. While modeling the first has proceeded smoothly, modeling the second has been problematic - a useful fact in determining the strengths and limitations of the model and the theory it embodies. Finally, Section 8 addresses some of the larger issues raised by the body of the paper. In the light of the four cases presented, I discuss the prospects, strengths, and limitations of the model. Finally, I return to a discussion of why engaging in this kind of enterprise, difficult as it is, might actually be useful.

2. On the use of models in educational research: What do we mean by a model, what properties do we want it to have, and what are the reasons for building one?

This section is intended to elaborate on what I mean by a model and to make clear both my expectations for models and their limitations. To discuss a series of issues regarding models and the modeling process, I shall consider various families of models. The first family of examples raises issues of description, explanation, and prediction. Consider the following picture, familiar from our school days:

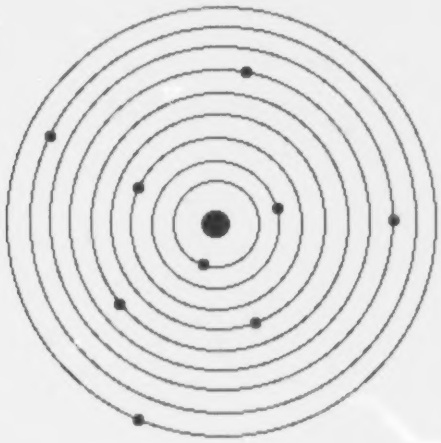

Figure 1
A 2-D "model" of the solar system

The two-dimensional representation in Figure 1, and perhaps its three- dimensional cousin (which has spheres of different dimensions held in different orbits by wires), are each called a "model of the solar system." In some sense they are models, of course - but they are very crude ones. Figure 1 presents some pictorial information. With supplementary text we come to understand that "the sun" is at the center of our solar system, and that nine planets orbit around it. Of course, this representation is grossly inadequate. The planets are not the same size; they are not co-planar; their orbits are not circular. While the 3-D physical model moves closer toward representing "reality" in that it shows non-co-planar orbits and indicates that the planets are of different sizes, those representations are more suggestive than accurate: the dimensions are not to scale, for example.

Now, one can imagine supplementing these models with a substantial amount of information. For example, one could give reasonably accurate descriptions of the

mass and of the orbital paths of each of the planets. Such specifications would provide considerably more accurate descriptions of the solar system than the other models - indeed, they would support testable predictions about the future motion of the planets. But even so, the model would be fundamentally lacking, in that the descriptions entail no sense of *mechanism* - no justification for why things are the way they are or will be. In the case of the solar system, that mechanism is provided by theories of gravitation. With the underlying theory, mere description is supplanted by explanation, and prediction is rationalized rather than just grounded in patterns. The kinds of models we seek are explanatory in the sense just described. They embody a theory of objects and relations among them.

The next set of questions I wish to address concerns the criteria by which one judges a set of theoretical claims or models. Three key criteria we employ are explanatory power, predictive power, and scope. I begin the exploration of this issue by considering a second family of models. Figure 2 offers my clumsy attempt at a representation of a helium atom.

Figure 2
A "model" of a helium atom

As in the case of the solar system (cf. Figure 1), this representation is woefully inadequate: one gets a suggestion of some electrons in orbit around a nucleus, but not much more. Also as in the case of the solar system, the story suggested by the picture can be expanded into a model that contains objects and relations among them, and that reflects an underlying theoretical structure. That theory, which lies at the foundations of modern chemistry, includes notions of atomic composition (an atomic nucleus is composed of protons and neutrons; the nucleus is surrounded by "shells" of electrons with a very particular structure), atomic weight, valence structure, and so on. As we know, that model works extraordinarily well. Consider, for example, its predictive power. The periodic table of elements was constructed as an explanatory structure, given then-current understandings of the chemical and physical properties of the known elements. Its structure reflected hypothesized regularities: the patterns of protons, neutrons, and electrons in the elements were predicted to grow predictably in certain ways, with the properties of the elements reflecting their atomic structures. When the known elements were placed in the periodic table, they fit - but there were gaps, suggesting that some elements (whose specific properties were predicted by their position in the periodic table) had not yet been discovered. The subsequent discovery of those elements provided strong

BEST COPY AVAILABLE

evidence that some of the assumptions underlying models of atomic structure were correct.

During the 19th and 20th centuries, the model was refined. It continued to provide the grounds for solid predictions. It suggested *why* certain atoms would combine, and in what ways - thus providing significant explanatory as well as predictive power. The model dealt with a wide range of phenomena - the structure of individual atoms, the ways in which they could or could not combine with atoms of other elements, etc. While the basic structure of the model remained more or less constant, descriptions of atomic structure evolved as scientists discovered more information. Over a century and more the field produced sequential refinements of the model, which, over time, seemed to be increasingly accurate and precise. It had predictive power, explanatory power, and tremendous scope.

Of great interest and relevance for the teacher-modeling enterprise is the fact that the atomic model described just above is likely to be wrong in some fundamental ways. Physicists no longer believe that atoms have shells of electrons with orbits as suggested in Figure 2: Heisenberg's uncertainty principle suggests that electrons cannot be identified as masses moving along particular trajectories. Physicists no longer speak in terms of the trajectories of individual electrons; they now talk in terms of "electron shrouds" and "activation states." The importance of this shift is as follows. Despite having all the wonderful qualities described above, the earlier models probably weren't "right" and their underlying theoretical foundations probably didn't describe the "truth." Those models have been supplanted by others, which may yet be supplanted by others. At the same time, it would miss the point to describe the early models of the atom as simply being "wrong." Whether or not they were ultimately correct, they were extraordinarily powerful in enabling us to understand and predict a wide range of phenomena.

In short, a model and the theory it embodies do not represent absolute claims to truth; they are working descriptions that help us grapple with complex phenomena. The models we build should be constantly tested against reality. They should judged by the ways they support prediction, by their explanatory power, and by their scope.[1] And, no matter how good they may seem, we should be prepared to jettison them when better models come along.

All well and good, the reader might say, but those are models of physical phenomena. Have I descended into the depths of scientism, and am I about to claim that such models capture human behavior? In short, my answers are "no" and "it depends on what you mean" respectively. All models are *representations* .

---

[1]By scope I mean the range of cases to which a model applies. A good model of the teaching process should apply to a wide range of teachers, for example, not just those with a particular style.

They are useful to the degree that they describe properties of the represented objects and help us to understand them. The meta-game of science is understanding the limitations of the representations one uses. At this point, any researcher would have to be deluded to claim that a model of human cognition or of human action-in-context "captures" a person to the point where the full complexity of that person's decision-making is accurately reflected in the representation. In any model of human behavior, things are clearly over- simplified. There is much that goes unrepresented. Some of the entities represented in the model (e.g., "beliefs," "schemata," etc.) may or may not have precise counterparts in the people being modeled. But, we may yet be able to put together a model that, in some ways, reflects the behaviors of teachers in action. And then (to deform an old saying): If it looks like a duck, if it walks like a duck, and if it quacks like a duck . . . then looked at in the proper way, it may help us to understand ducks a bit better.

A main goal of the Teacher Model Group's research is to use the model-building process to construct rich portrayals of the ways people interact with the world, to test the accuracy of those portrayals, and to be appropriately cautious in spelling out the implications for people of what is seen in the model. Our goal is not to create little robot teachers, or computer programs that act like particular teachers. Rather, the goal is to understand what teachers do in the classroom - how and why they make and implement decisions in the course of teaching. Like all of the models described in this section, our model must have two main parts: components and mechanism (that is, we must describe the parts of the model and how they work together). As suggested in Section 1 and as will be elaborated in Section 3, the main components of the model are the teacher's beliefs, goals, and knowledge. I shall, in Section 3 and in the case studies, devote a fair amount of time to discussions of how those components interact. As the reader works through the description and the case studies, he or she might want to keep the following criteria in mind: How well do the theory and the model "explain" what the teacher does? (Is there a solidly worked out sense of mechanism? Is it believable? Is there enough evidence to suggest the adequacy of the warrants for the claims being made?) Could one envision the model and the theory used for purposes of prediction? (That is, if you had developed a model of a teacher teaching in a particular context, would that model make reasonably accurate predictions about what the teacher would do in a closely related context?) And, what is the likely scope of the theory? (How large is the range of teaching styles and experiences that one could "capture" effectively using the model?)

3. ELEMENTS OF A THEORY.

Figure 3 provides a crude pictorial representation of the major components of our model of teaching-in-context. Like the sketches in Figures 1 and 2, Figure 3 is meant to be suggestive rather than directly representative. The text that follows elaborates on the components of the model and the relationships among them.

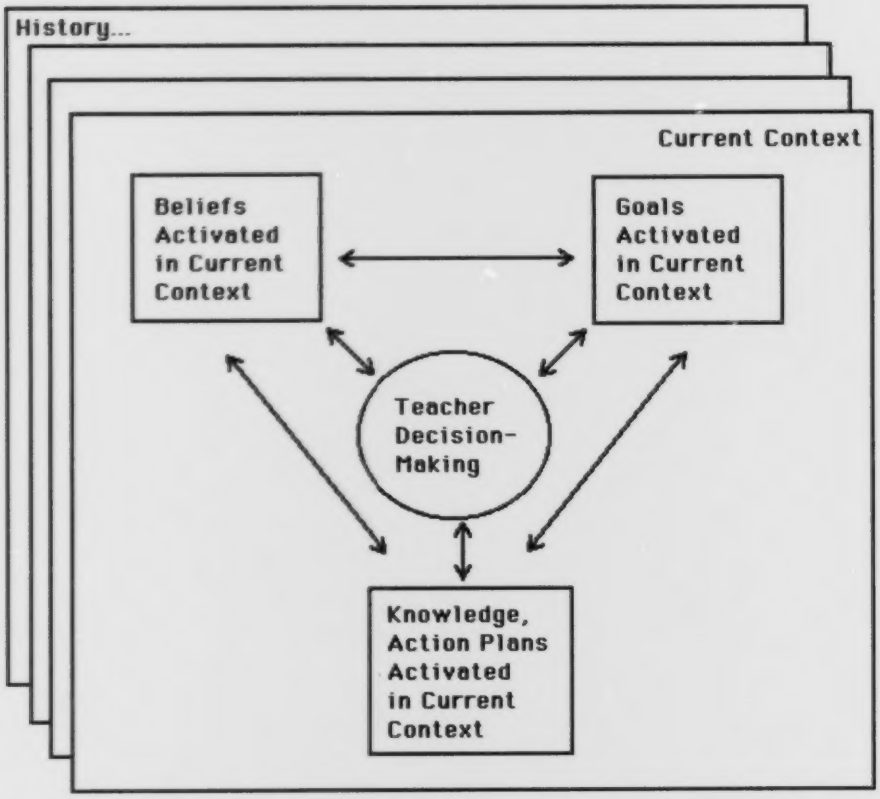

Figure 3.
Aspects of a model of teaching-in-context

The discussion of Figure 3 that follows takes place in two layers of unfolding detail. Section 3.1 provides a general description of the major components of the model and how they fit together. This is followed in Section 3.2 by a more

complete description of each of those components and of the theoretical mechanisms that link them.

3.1. What's in the model, and how does it work? A top-level view.

At the most cursory level, Figure 3 can be seen as representing the following.

* Teaching takes place in, and is a function of, the instructional context. Teaching is a dynamic act, responsive to what happens in interaction with the students. The teacher is constantly monitoring what is taking place during instruction and acting on the basis of perceptions of what is taking place.[2]

* What takes place in the current context is, of course, a function of history - the teacher's history, the students' histories, what has happened in that classroom between teachers and students, etc.

* The three boxes in the front plane of Figure 3 represent major components of the model: beliefs, goals, and knowledge. None of these necessarily have priority over any of the others - as elaborated below, they all affect each other. That is, one might start with beliefs, asking how the teacher's beliefs shape the teacher's goals, and how those in turn shape the actions taken by the teacher. But any other pathway through Figure 3 gives rise to equally important questions. Note that the boxes representing beliefs, goals, and knowledge are labeled with the phrase "activated in current context." Those labels are there to stress what is happening *at the moment*. That is, a teacher may have numerous beliefs that shape his or her teaching through the year. However, at any particular time, some of those beliefs are strongly activated, either because of prior planning or because their activation level was influenced by an event that has recently taken place; others are inactive or activated at low levels. Similar comments hold for goals and knowledge.

What follows is a slight elaboration on the contents of the three boxes.

* The knowledge base. What a teacher might do in any situation is, of course, fundamentally shaped by the set of intellectual resources the teacher can bring to that situation - that is, the teacher's knowledge base. This large category includes knowledge of the students, of the context, and of the content. It includes a variety of general and content-specific classroom and interactive routines. Central among these are "action plans" at various levels of grain size that can be used to achieve various goals (e.g., the expectation of using a standard routine for collecting homework, an interactive dialogue intended to engage students with a particular idea, or a well rehearsed "mini-lecture" on a particular topic). The kinds of resources the teacher can bring to bear, and how

[2]Of course, the degree of monitoring and responsiveness can vary greatly. There is little of either in the "classic lecture," and a great deal in highly interactive classrooms.

such resources are accessed, are major issues for modeling. See Section 3.2 for detail.

* Goals. What a teacher does at any moment is shaped by the teacher's goals. These range from very long-term ("I want students to understand mathematics as a sense-making activity"; "I want them to become independent learners") to medium-term ("the students need to understand derivative as rate of change"; "they need to collaborate well on this activity") to short-term ("the solution to this exercise should be on the board clearly for their notes"; "I want it quiet *now* ") More precisely, decisions about what to do and how to do it are shaped by the set of *currently active, high priority* goals. Suppose that a teacher wants to achieve X. The teacher might have access to various action plans that could result in X. Which of those action plans is actually selected will depend on a number of things, including the constellation of active goals (one action plan might help achieve X and Y, while another may help achieve X and Z; how important is Y relative to Z?) and perceived constraints (Suppose a particular approach has "value added" in that it helps to achieve X and Y, as opposed to "just" X. However, it takes much longer to implement. Can time for this approach be made? If so, what gets sacrificed?).

* Beliefs. Which goals and action plans have high priority (at any particular time, in any particular context) is a function of the teacher's beliefs and epistemology. What is the teacher's sense of the mathematical enterprise? What counts for "understanding"? How important are (for example) formulas, explanations, qualitative reasoning? What does the teacher think these students are capable of learning? What counts for proof, for this audience? How important is discipline? (Is a good classroom a quiet classroom? How does that play out when the teacher chooses between individual or group work for the class to engage in?) How the teacher feels about all of these issues, consciously or unconsciously, will be a major factor in determining which goals have highest priority. These, in turn, will shape what the teacher chooses to do.

* Mechanism. The issue is, how and why does the teacher do what he or she does at any moment? By way of preliminaries, I note that the structure of the model is self-similar, in that the same basic mechanisms work at multiple levels of detail. That is, the descriptions that follow are appropriate for thinking about the whole year's instruction, about a particular unit, about a particular lesson or lesson segment, about the discussion of a particular problem, or about a brief exchange with a student or the class.

The simple version of the story is this. At any given moment, there is a constellation of highly activated beliefs, goals, and knowledge. If things are going well, these are all nicely in synch. The beliefs serve as a backdrop for what the teacher is trying to accomplish - the goals. The teacher has accessed various kinds of knowledge in the service of those goals, and has a set of expectations about the ways things are likely to unfold. (This set of expectations, referred to

as a "lesson image," will be elaborated more fully in Section 3.2.) In concert with these expectations, the teacher is engaging in some set of activities with the students (that is, carrying out an action plan consistent with the goals and beliefs).

If things go according to plan,[3] the teacher's lesson image and action plan are the predominant factors in shaping what happens next. For example, routine actions such as collecting homework, calling on students, or working through an example on the board will, barring untoward events, take place pretty much as envisioned; likewise for transitions from one event to another, such as the transition from reviewing homework to introducing a new topic. As such transitions are made, there is a dynamic reprioritization of clusters of beliefs, goals, and knowledge. Goals that have been satisfied lose their high priority status, and new goals consonant with the new activities replace them. Similarly, concomitant with the engagement in the new activities is the activation of relevant beliefs and knowledge.[4]

Where life gets interesting, of course (and it does all the time!) is when things do not proceed according to plan. (Note that this may happen at every level of grain size from the trivial and instantly fixable, where a quick action may set things back on track, to the substantial, where an issue arises that, if addressed, will take a substantial amount of time to sort out.) Here is a generic description of what happens. The baseline is the current state - the currently activated set of beliefs, goals, and knowledge of the teacher, and the current set of activities in which the class is engaged. Something happens. Depending on its nature, and the context, this event may trigger the high activation of some currently low-activation beliefs.[5] In the context of these beliefs, the given constraints (e.g., Is there enough

---

[3]The phrase "according to plan" needs elaboration. Typically, plans for teaching are rough approximations, with much to be filled in at the moment; hence "according to plan" is a relative judgment. By analogy, consider a drive I might plan to make from point A to point B. In action, my plan unfolds at multiple levels of detail. There are major decisions such as which roads to take, and then, once I am on a particular road, much more local decisions - making lane changes, slowing down to avoid an aggressive driver, etc. Do things proceed "according to plan" when I swerve to avoid a pothole, move into an empty lane, get off an exit early to avoid traffic congestion, or decide to take a different route because of a traffic report on the radio? How you answer depends in part on which plan you focus on - the local one to continue driving safely down this road right now, or the global one to get from point A to point B.

[4]I should note that our choice of language, which is consistent with the kind of goal-driven architectures of AI models, may tend to obscure the similarity between our work and that of others. For example, Leinhardt and colleagues (Leinhardt, 1993; Leinhardt & Greeno, 1986; Leinhardt, Weidman, & Hammond, 1987) use the terms *agendas* and *routines* in ways that parallel our use of *lesson images* and *action plans* respectively. There are strong parallels between our conceptualization of "how teaching works" and theirs.

[5]For example, a disrespectful comment made by one student about another may activate the teacher's beliefs about appropriate classroom norms. Or, a question by a student may raise interesting issues but will take a lot of time to pursue. How important does the teacher believe it is to "honor student inquiry"? Or, a comment by a student may reveal a misunderstanding about X. How important is to raise X with the whole class?

time to deal with this?) and resources (What knowledge gets activated? What action plans are available? Will they provide a way to deal with this issue, or will the teacher need to march off into uncharted territory?), new high-priority goals may be formulated and a new action plan consistent with the new constellation of high-activation beliefs and goals may be put in place.

In summary: The model is context- and history-sensitive, in that the teacher's decisions and choice of actions are responsive to the immediate context, the teacher's personal history, the teacher's history with the students, and the teacher's "active" (high priority) beliefs, goals, and action plans. Classroom occurrences are perceived and interpreted in the light of the teacher's beliefs, and the teacher's decisions are made in the light of high priority goals (which are shaped by beliefs and accessible resources) and mechanisms for achieving them (the teacher's knowledge, which includes access to a collection of potential action plans). At any point, things may proceed according to current expectations, in which case the teacher continues with the implementation of the current action plan. Or, new events may result in a different activation pattern of beliefs, knowledge, and goals, resulting in the selection and implementation of a new action plan. As in problem solving of any sort, the decision-making may or may not be conscious, and it may or may not involve reflection on possibilities and consequences.

3.2. Elaboration: Definitions and more detail

This section of the paper offers a brief annotated glossary of major terms used in this paper. Without further ado:

Theory and model

Given the extensive discussion in Section 2, little needs to be said about these two terms. Both are used in the general scientific sense. Theories should explain things - that is, they should work at a level of mechanism. Models are the embodiments of theories, which represent the entities described in the theories and the relationships among them.

It may be worth saying explicitly what I do not mean by the term "model." I do not mean model in the sense of exemplary: "here is a model teacher." As stressed in Section 1.2, a good theory should allow us to model teaching of all types. Moreover, this work is not prescriptive, in the sense of providing a model of "how to do it." Of course, insights gained from modeling proficient teachers can be valuable in that regard. But using such insights as a guide to improve teaching is an application of the modeling work described here, and should not be confused with the modeling enterprise itself.

Theory of Teaching in Context Draft EE Alan Schoenfeld 7/7/98

Lesson image

A teacher's lesson image is, in a very expansive sense, the teacher's envisioning of the possibilities and contingencies related to a lesson. The teacher's lesson image includes knowledge of his or her students and how they may react to parts of the planned lesson; it includes a sense of what students are likely to be confused about, and how the teacher might deal with that confusion; and more. As such, the concept of lesson image is related to but not subsumed by the concept of teacher planning. For an early, extensive review of the literature on teacher planning, see Clark and Peterson (1986); for more recent perspectives see Leinhardt (1993) and Calderhead (1996). The term "lesson image" was introduced by Morine-Dershimer (1978-79).[6] It is of particular interest to us given our wish to understand teaching-in-context - after all, teachers' envisionings of what they expect to take place in the classroom play a major role in shaping what does take place. Teachers' knowledge of instruction, captured partly in their images of what it will look like, has been described in various ways: see, e.g., Elbaz (1983) and Clandinin (1986). In work that is closely related to ours, Leinhardt (1993, pp. 22-23) writes about teachers' *agendas*. Leinhardt describes agendas as follows:

1. They include actions of both the teachers and the students.
2. They refer to some prediction of student behavior.
3. They frequently suggest or make reference to a test or a check that helps determine how to proceed.
4. They mention the location of the particular lesson in the wider spectrum of lessons.
5. They often include such overarching pedagogical rules such as moving from the concrete to the abstract or such overarching, context-driven rules as "this idea is useful in understanding the next idea."

To illustrate the notion of lesson image I shall provide a specific example, my lesson image for the opening day of my undergraduate course in mathematical problem solving. In addition to illustrating the concept of lesson image, this discussion serves as a backdrop for a brief analysis of the lesson itself. That analysis, which is given in Section 7.1, shows how the lesson image described here actually played out during instruction.

The opening days of my problem solving course are crucial, in that they set the tone and direction for the entire course. I must convey the spirit of the enterprise; I must begin to create the classroom climate that will support productive mathematical interactions; I must convince the students that they

[6] A point Morine-Dershimer (1979-79) stressed is that a lesson plan, the written artifact that captures the planned structure of a lesson, may represent only a small fraction of the "shaping structure" with which the teacher enters the classroom.

have something of value to learn in the course. Having developed the course more than twenty years ago and having reflected on it frequently since then, I have a very strong sense of what I want to have happen at the beginning of the course, and of the things that I will do to ensure that it will happen. What follows is a general description thereof.

The first day will begin with a fifteen-minute introduction to the course. The introduction provides some background and describes the basic course structure. Having honed the introductory lecture through the years, I can do it pretty much seat-of-the-pants; a quick review of main points before class insures that I will cover the things I want to, and remember to tell the specific motivational anecdotes I want to tell. When the introduction is done I will hand out a sheet of problems, asking the students to break into small groups and work together on the problems. As they work on the problems I will circulate through the class, facilitating the discussions and observing the things they do while working on the problems. Facilitating discussions can be done impromptu. It is a well- practiced skill, and the problems have been chosen carefully to support productive interactions among the students. After about twenty minutes of small-group discussions, I will call the class together for whole-group discussions. We will start to work the problems on the assignment sheet.

The discussion of the first problem is carefully staged. The problem is one I know the students will have difficulty with, although they have actually learned to solve it in previous classes. This provides the opportunity for a "mock lecture" in which I will caricature some of the students' previous experiences in formal mathematics courses and set them up for a different kind of experience in this one. We will then move on to the other problems.

On the one hand, the whole-class discussion that takes place from this point on is rather free-form: I will ask the students what they came up with, and I will react to it. From the students' point of view, the discussion will appear impromptu - and in a sense it is. On the other hand, having taught the course for many years and having used these problems in a large number of workshops, I know exactly what to expect for the discussion of every single problem. I can tell you what questions I will ask, what the students will respond (collectively, not individually) when I ask those questions, and how I will respond to what they say. Of course, people will not always say the same things, and not in the same order - but in a general sense, I know what to expect. The structure of the discussions is largely constraint-based, in that within instructional segments students are asked for comments and issues are usually discussed in the order in which they arise, rather than in a pre-determined order.

In short, I can tell you, before the class starts, how things are likely to unfold - down to an extremely fine level of detail. It is not that I follow a rigid plan and coerce the students into it; there are many branch points and contingencies.

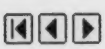

However, I know what most of them are likely to be. And, there are few surprises.[7]

Now, this example of a lesson image is an extreme case. I have a lot invested in making the first days of the course work in certain ways, and a lot of experience in teaching the course. Nonetheless, it makes the point that I walk into the classroom with a lot more than my lesson plan. (In fact, the only written artifacts I use for preparation are my first day's handout describing the course and giving the set of problems I intend to discuss.) The same is the case in general. Lesson images vary widely from teacher to teacher, and from context to context for the same teacher. They may be tightly scripted, or they may be very loose. Either way, these envisionings of what will take place play a major role in shaping what will take place. In modeling any particular teacher, then it is important to understand the teacher's lesson image for the instruction being modeled.

Beliefs

Beliefs are mental constructs that represent the codifications of people's experiences and understandings. Teachers have beliefs about themselves (e.g., they might believe they are good or bad at mathematics), the nature of intellectual ability (some people believe it to be innate, others that it is malleable), about the nature of the discipline they teach, about learning, about individual students, about groups of students, about the environment in which they work, and more. People's beliefs shape what they perceive in any set of circumstances, what they consider to be possible or appropriate in those circumstances, the goals they might establish in those circumstances, and the knowledge they might bring to bear in them.

There is a huge literature on mathematics-related beliefs in general, and teachers' beliefs in particular (see, e.g., Aguirre & Speer, 1996; Borko & Putnam, 1996; Calderhead, 1996; Cohen, 1990; Cooney, 1985; Ernest, 1989; Lampert, 1990; Pajares, 1992; Schoenfeld, 1985, 1992; Strauss & Shiloney, 1994; Thompson, 1992). In what follows I shall briefly sketch out a general description of beliefs, and then some relevant classes of beliefs for the model.

I take two main points from the literature:

---

[7]When there is a surprise, it is dealt with in the manner suggested above. Sometimes my reaction is ad hoc; sometimes what occurs results in my doing something unexpected (in this context) but familiar. The next time I teach the course, the possibility of this event's occurrence and my possible reactions to it are incorporated into my lesson image. For example, one of the first times I taught the course I discovered when using a particular problem that some students were weak on mathematical induction, which was one of the methods used to solve it. That time I did an impromptu mini-lesson on induction - easy enough, because induction is a standard topic. The next time I taught the course, I was prepared for a detour into induction, with some better-chosen examples and a better explanation.

1. Beliefs have a strong shaping effect on behavior.
2. There is a major difference between *professed* beliefs and *attributed* beliefs. In our modeling efforts, we use professed beliefs as one of many sources of evidence; but the model of that individual's teaching contains beliefs we attribute to the individual.

Point 1 above has been amply documented (see the references cited above), so I shall assert it here without further amplification. Point 2 above is essential to understand, and merits some elaboration. There are two key issues here. The first issue is that people may say A and do B, and the two may not be compatible. Hence, we must distinguish between their *professed* beliefs and the beliefs that underlie actual behavior. When people behave in certain ways, we *attribute* beliefs to them. As noted, these attributions may or may not correspond to those people's professed beliefs. The second issue is that I want to be as precise as possible about the attribution process. We can never know what someone truly believes. Hence, when we attribute beliefs to someone (or to a model of that person's behavior), what we are really saying is: "this person behaves in a way that is consistent with his or her having those beliefs." I will use the shorthand, but the caveat should be understood.[8]

An expanding literature on teacher beliefs points to the subtleties and complexities concerning them. The relationship between belief and action is far from straightforward - indeed, relationships among beliefs are far from straightforward. As studies like those by Cooney (1985), Cohen (1990), and Aguirre and Speer (1996) indicate, beliefs can (a) be contradictory, and (b) have indirect but strong effects on teaching practice. The subject of Cooney's investigation had a heartfelt and strongly professed belief in the value of "problem solving," but his classes showed little evidence of Pólya-like influence. The reason? The teacher's understanding of problem solving (largely as recreation, as opposed to as a deep way of conceiving mathematics) led to his trying to use "interesting" problems largely for motivational and recreational purposes. His students neither enjoyed nor saw the value of such problems, and the teacher found his interest in problem solving outweighed by the need to make sure his students "got" the regular material. Cohen's study indicates that a teacher can believe that he or she is teaching in the spirit of reform while at the same time employing instructional methods that reflect beliefs about mathematical teaching and learning contrary to those that underlie reform efforts. Likewise, a fine-grained analysis by Aguirre and Speer (1996) indicates that a "reform-oriented" teacher who professed a belief in having students contribute substantially to content discussions put a quick end to such contributions when they did not result in what she wanted aired in the classroom - this contradicted another, stronger belief she had, which is that what makes its way onto the board, and thus the students' notes, should be crisp and correct.

---

[8]For a general discussion of this issue see Pajares, 1992, esp. pages 316, 326.

Theory of Teaching in Context Draft EE Alan Schoenfeld 7/7/98

It is clear that the following classes of beliefs affect teachers' classroom actions, and must be examined in a comprehensive model of teaching:

* beliefs about the nature of subject matter (in general and with regard to the specific topics being taught);
* beliefs about the nature of the learning process (both cognitive and affective);
* beliefs about the nature of the teaching process and the roles of various kinds of instruction;
* beliefs about particular students and classes of students.

The research literature makes it clear that beliefs are often context-dependent, that they have differing strengths in differing contexts, that they tend to be activated in clusters (one highly activated belief may trigger the activation of other closely related beliefs), and that conflicting beliefs can "compete" for priority. When beliefs are represented in a model of teaching-in-context, they should, likewise, have these properties.

Goals

A goal is something you want to accomplish. Central to this discussion of goals are the following five issues.

First, I should be careful to make the same kinds of distinctions with regard to attributed versus professed goals that were made above with regard to attributed versus professed beliefs. Some goals may be explicit, some tacit and unarticulated. Teachers may profess certain goals, and yet their actions may undermine them (possibly because conflicting goals are also operative, possibly because there can be straightforward differences between what one professes and what one does). Most importantly, while modeling a specific teacher we are *attributing* particular goals to that teacher. Of course, that attribution is done carefully, and with as much triangulation as we can provide. Nonetheless, the careful reading of any goal attribution should be: "the teacher is behaving in a way that is consistent with his or her having those goals (at this time, with the level of priority indicated)."

Second, goals occur at many grain sizes and can have high activation levels for widely varied amounts of time. Some goals, which are called *overarching goals* , may be present at some level during all of instruction, rising to high activation status periodically. Examples of goals of this type are "I want my students to experience mathematics as a sense-making activity" and "I want this classroom to function as a respectful and collaborative intellectual community." These goals may play different roles at different times. At some times they may have highest priority - e.g., when the teacher is trying to shape the character of the classroom community, or having the students engage in particular activities designed to

highlight the sense-making aspects of the discipline (see the discussion of Minstrell's benchmark lesson in Section 6). Sometimes overarching goals play a shaping role in the choice of activities. For example, action plans A and B might be two different ways of achieving a current high priority goal. If B also contributes to the achievement of an overarching goal while A does not, then B may be chosen. And sometimes these goals seem to reside in the background, with no current effect. However, if an event occurs that contradicts them (e.g., a student makes a comment that has the potential to disrupt the sense of classroom community the teacher is trying to build or maintain) then the goals are suddenly activated with high priority. Goals range from long-term (extending over days and weeks) to short-term ("I want to get through preliminaries fast, because we have a lot to achieve today").

Third, goals may be pre-determined, or they may be *emergent*. When things go according to plan (i.e., lesson image), new goals click into place as current goals are satisfied. For example, having finished my introductory lecture in the problem solving course and handed out my problem set, my goals switch to facilitating small-group discussions and monitoring student progress. When unanticipated events take place - e.g., an activity doesn't go as planned, or a student makes a suggestion that warrants following up - then new (or newly prioritized) goals become activated.

Fourth, it should be stressed that there is no simple correspondence between goals and actions - there are not necessarily simple linkages of the type "If the teacher wants to achieve X, the teacher does Y." (Such connections can happen, of course, but they are relatively rare.) Rather, at any time, the teacher is likely to have a constellation of active goals, at various levels of activation. This constellation of goals shapes the actions taken by the teacher, in that at any given time the action plan the teacher chooses to implement will be selected on the basis of how well it satisfies the high priority goals and constraints currently in place.

Fifth and finally, I want to re-emphasize a point made in Section 3.1, that goals and actions can be co-emergent. As suggested in the previous paragraph, the most straightforward (and modelable) explanation of behavior is that actions are taken in the service of goals. But, reality does not always correspond to this assumption. Sometimes one responds almost automatically to an occurrence in the classroom. Such actions are clearly not taken to satisfy a pre-existing goal. If, however, the action taken by the teacher clearly satisfies some goal in response to the situation, I will want (for purposes of modeling) to attribute that goal to the teacher.

The knowledge base

One could write volumes about the knowledge base, both in general and specifically regarding teacher cognition. For extensive reviews, see Borko and

Putnam (1996), Calderhead (1996), Clark and Peterson (1986), Fennema and Franke (1992), and Fenstermacher (1994); for a perspective on alternative descriptions of teacher knowledge, see Sherin, Sherin, and Madanes (1996).

Here I discuss in brief two ways of thinking about teacher knowledge: the knowledge "inventory" (what one knows), and how that knowledge is organized and accessed. A simple metaphor about libraries explains the difference between the two views. The first kind of question one thinks to ask about a library is typically related to its contents: How good is the library's collection? Does it have a good reference section on teaching and learning (or contemporary novels, or 15th century Spanish literature, or periodicals, or . . .)? Content questions such as these can be answered solely with the library's catalogue - one need not know anything about the physical layout of the library, how and where things are shelved, etc. A second kind of question concerns getting access to the contents: If I am interested in a particular book or journal, how do I find it? How do I find related works? In the case of these questions, the nature of the storage system is of central concern. It is worth noting that when one focuses on content questions like those mentioned above, "process" questions on the order of "How do I access X?" are essentially irrelevant. Conversely, process questions may be essentially content-independent: I may follow exactly the same procedure to locate a book on the psychology of mathematical cognition that I would follow to find a book on 15th century Spanish literature.

The same is the case with regard to knowledge. Descriptions of various kinds of knowledge (subject matter knowledge, pedagogical knowledge, practical knowledge, etc.) may result in the view that there are various, separate categories of knowledge. However, this separation into categories may be deceptive. As in the case of the library, there may be hundreds of different categories - but only a small number of fundamental ways in which the knowledge in those categories is accessed for use. For example, the way I access my mathematical content knowledge may be fundamentally the same as the way in which I access my general pedagogical knowledge.

The knowledge inventory.

I begin with a brief description of the knowledge inventory. Borko and Putnam's 1996 review sets out a mainstream framework for considering the content of teacher knowledge and beliefs. Here I focus on the knowledge aspect. Following Shulman (1986), Borko and Putnam identify three main categories of knowledge: general pedagogical knowledge, subject matter knowledge, and pedagogical content knowledge. Brief descriptions and commentaries are as follows:

General pedagogical knowledge "encompasses a teacher's knowledge . . . about teaching, learning, and learners that transcend particular subject matter domains. It includes . . . classroom management, instructional strategies for conducting lessons and creating learning environments, and more fundamental knowledge

and beliefs about learners, how they learn, and how that learning can be fostered by teaching." (p. 675)

Subject matter knowledge : "What is essential to recognize is the argument that teachers need to know more than just the facts, terms, and concepts of a discipline. Their knowledge of the organizing ideas, connections among ideas, ways of thinking and arguing, and knowledge growth within the discipline is an important factor in how they will teach the subject." (p. 676)

Pedagogical content knowledge. Borko and Putnam focus on Grossman's (1990) elaboration of the concept, which contains four major categories of components:

(1) "the teacher's overarching conception of the purposes for teaching a subject matter . . . the nature of the subject and what is important for students to learn";

(2) "knowledge of students' understandings and potential misunderstandings of a subject area . . . [including] preconceptions, misconceptions, and alternative conceptions about topics such as division of fractions, negative numbers. . . . ";

(3) "knowledge of curriculum and curricular materials"; and

(4) "knowledge of strategies and representations for teaching particular topics." (pp. 676-677)

Noting that the decomposition given is just one of many possible ways to delineate knowledge, Borko and Putnam (1996, p. 677) point to alternative frames: "narrative forms of knowing (Bruner, 1986), . . . situated knowledge (Leinhardt, 1988), event-structured knowledge (Carter & Doyle, 1987), personal practical knowledge (Connelly & Clandinin, 1985; Elbaz, 1983), images (Calderhead, 1988; Clandinin, 1986), and knowledge in action (Schön, 1982)." To these one might add these categories from Calderhead's (1996) classification: craft knowledge, case knowledge, theoretical knowledge, metaphors. To return to the library metaphor, however: these manifold, different, and overlapping categories of knowledge represent the contents of knowledge. A second issue, which may cut across many of these categories, is how that knowledge is accessed and used. That issue, pursued immediately below, is central for a theory of teaching-in- context.

The organization and access of knowledge for use in teaching.

It is widely accepted in the psychological literature that people organize their experiences mentally via mental representations of familiar *classes* of experience. The most common name for such abstractions is *schema* (plural *schemata*); related terms are scripts and frames. Attached to a schema are its typical features, some knowledge related to it, and typical ways of behaving when that schema has been called to mind. The classic example demonstrating the existence of such knowledge structures in mathematics learning is that of a schema for solving

word problems in algebra. Hinsley, Hayes, and Simon (1977, p. 97) documented that students, when given just the first few words of a word problem such as "a river boat is . . . " will consistently recognize the class of problems it came from and say things like: "It's going to be one of those river things with upstream, downstream, and still water. You are going to compare times upstream and downstream - if the time is constant, it'll be distance." That's a lot to read into a few words, but we do that kind of thing all the time - not just with regard to words, but with regard to almost every kind of phenomenon. For example, the p-prims identified by diSessa (1993) are a similar type of abstraction of experience.

This is a general point about human behavior, not constrained to mathematics teaching or learning. The core idea is that we humans abstract our experiences in the world, and we use those abstractions as means of perceiving and interpreting things as we interact with them. The salient features of an object or a situation may cause us to identify that object or situation as an X. Once we have done so, we expect it to behave like an X: we look for other typical features of X's, predict its properties and behavior according to the typical properties and behavior of X's, and so on. With extended experience much of this becomes automatic - we see things in certain ways, and "know" their properties. In turn, this kind of automaticity frees up cognitive resources and supports flexibility. In addition, the more familiar we are with a domain, the more fine-tuned our perceptions and schemata.

Specifically with regard to teaching, it can be argued that much of teacher expertise can be seen as the result of the development of such abstractions. Shavelson (1986) refers to a range of schema types that contribute to accomplished teaching, for example the *scene schema*, which represents teachers' "ability to rapidly recognize [and then act on] common activity structures that, to a novice, might appear as chaos" (Shavelson, 1986, p. 4). Berliner (1994) posits that "experts often develop automaticity for the repetitive operations that are needed to accomplish their goals" (p. 169), "experts are more sensitive to task demands and social situation when solving problems" (p. 171); "experts are more opportunistic and flexible in their teaching than are novices" (p. 173); "Experts have fast and accurate pattern recognition capabilities" (p 177).

One main feature of human memory, then, is that elements in memory are organized in "chunks." In what follows I shall delineate some of the most important kinds of chunks related to teaching: action plans, routines, mini- scripts, and more. A second main feature of human memory is that it is *associative:* when a particular item "comes to mind," related items are likely to do so as well. For example, if you mention "Pythagoras" to me in almost any context, I think of the Pythagorean theorem. I am likely to envision a picture of a right triangle that has sides labeled a and b and a hypotenuse labeled c, along with of the equation $a^2 + b^2 = c^2$. Examples such as the (3,4,5) and (5,12,13) right triangles may spring to mind, and, depending on the context, I may think of the

general solution in integers to the Diophantine equation $a^2 + b^2 = c^2$. If the topic comes up in a classroom context, I may also think about productive ways in which my students and I have worked on issues related to the Pythagorean theorem.

Two quick comments are appropriate here. First, although the example about the associativity of human memory just given was primarily mathematical and content-related, it should be stressed that the phenomenon is more general. All of memory is associative. For example, bringing a particular item to mind may call up not only related knowledge but related beliefs and emotions as well; an event that triggers the activation of particular beliefs may trigger the activation of knowledge related to those beliefs, and so on. Second, it should be noted that there are rigorous ways (e.g., spreading activation networks) to model the kinds of associative structures that have been described informally in this paper. Phrases such as "the activation of X triggers the activation of Y" can be translated into specific detail in a computational model, where the strengths of links between objects modeled in memory are delineated, where changes in the activation level of one object are propagated through the network to produce changes in the activation levels of everything connected to that object, and where "X comes to mind" has the specific meaning that the activation level of X exceeds a particular pre-set value. Such detail is not discussed in this paper. The focus here is on human behavior, so computational issues will not be pursued.

Action plans.

A central construct in the model of teaching is the *action plan*. Roughly speaking, action plans are prospective means for achieving goals - that is, an action plan is a set of actions intended to be taken in order to work toward the achievement of a constellation of current high priority goals. I devote significant attention to these because the process of modeling specific teachers consists of understanding their action plans and the ways they are implemented.

Here I briefly indicate some of the properties of action plans and identify various types of action plans that are frequently used in teaching. The way I do so is to first provide some examples - I provide an overview of some of the action plans used in the first day of my problem solving course. With the examples in hand, I will describe some of the properties of action plans in general.

By way of context, recall the discussion in the section on "lesson image" of how I envision the opening sessions of my problem solving course. Specifically, I enter the classroom the first day with this general plan: I will (1) provide an extended introduction to the nature of the course and what students can expect to get out of it, (2) have students work on problems in small groups while I circulate through the classroom, (3) hold whole-class discussions of the first few problems

until the class period is almost over, and then wrap up the discussion with a particular problem that ends the class on a high note.

Each of these major chunks in the general plan has a corresponding action plan, which will itself unfold further in detail. Taking them in order:

(1) The approximately fifteen-minute introduction to the course is in essence a lecture-performance that follows a loose script. Components of that script include laying out specific information about the course and telling some specific anecdotes about problem solving in general and about past students' successes in the course. Some of these components may be decomposed further, but many are already at a level where there is no need for me to decompose them - I feel comfortable that I can tell particular stories when I need to.

(2) Having students work on problems in small groups while I circulate through the classroom is a routine that I introduce early, and use consistently throughout the course. Doing so serves two purposes: (a) I can facilitate small group interactions and make sure that the discussions are "on track"; (b) I can gather information about what the students have done, in preparation for the whole- class discussion to follow. The actions that support both of these purposes are familiar and done on an ad hoc basis. However, there is a structure to the ad hoc-ness. I have standard interactive routines for approaching groups of students and working with them, and when students say particular things those routines are likely to be invoked. (That is, the goals and the action plans to meet them are emergent, in response to what students say and do - but much of the support structure is available for immediate access in memory.)

(3) For each of the problems we discuss, I have a set of elaborately worked-out scenarios that include branch points in the conversation, ideas about what students are likely to say and how I will respond, and more - e.g., I am prepared to do a mini-lecture on induction in the middle of the discussion of one problem, because I know the need for it has arisen in the past (see Arcavi, Kessel, Meira, and Smith, 1998, for detail). The action plans for these discussions unfold in layers of detail. For example, I will ask how we might think about solutions to the problem. Three methods might be proposed, and unless there is reason to do otherwise, we will discuss the suggestions in the order that they have arisen. Based on my experience, the odds are that I can anticipate each of the suggestions; hence when they arise, I can employ my action plan for the discussion of that particular approach. Note that the discussions are generally constraint-based rather than strictly linear (that is, the order of approaches discussed is not pre-determined), but that some of those constraints can be linear: for example we tend to work through all the solutions to a given problem before we consider generalizations of it.

Theory of Teaching in Context Draft EE Alan Schoenfeld 7/7/98

Here are some properties of action plans. Action plans:

* occur at various levels of grain size. Just as goals occur at various levels of grain size (from overarching goals such as "I want students to experience mathematics as a sense-making activity" to short-term goals such as "I need to make sure that this definition is written clearly on the board and copied into the students' notebooks"), so do action plans.

* are often nested, unfolding in detail. (See the discussion of item 3 immediately above, where plans for a discussion contain plans for parts of the discussion, which contain . . .)

* often work toward achieving a constellation of goals - e.g., moving the discussion of a problem forward while helping to build a classroom community of inquiry.

* may come "pre-packaged," as part of a lesson image, or they may be implemented on an ad hoc basis in response to emergent goals in the classroom.

* may "spring from memory" in essentially full form, or may be created on the spot - e.g., when a student asks a question and the action plan created is "figure out the mathematics quickly and present the solution to the class."

* may be relatively complete, with activities scripted down to a relatively fine-grained level. But, action plans may also be quite incomplete - e.g., "I'll take questions for up to ten minutes and deal with whatever comes up."

Some of the typical forms in which action plans appear are:

* Routines. These are established patterns of behavior that often cut across the specifics of content - e.g., starting the day with a review, asking for and dealing with student comments and questions, collecting homework in particular ways, having students come to the board, and so on.

* Scripts. These are content-specific, imagined scenarios for the ways in which discussions will play out - e.g., my full envisionings of the discussions of the first few problems in my problem solving course. Note that scripts may be highly flexible and interactive, with "slots" for student actions, and with contingencies for responses depending on what students produce and when they produce it (or don't). I use the term "script" in the sense of the cognitive literature (see, e.g., Schank & Abelson's 1977 discussion of restaurant scripts).

* Mini-lectures are a special case of scripts which are somewhat screenplay-like in nature. Here the teacher is asked a question for which there is a

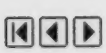

familiar, essentially "packaged" response, and the response is delivered directly.

* Simple talk. Sometimes the way one deals with a (small scale) situation is to simply produce a brief explanation or comment, formulating it on the spot and saying it as it is formulated. A large fraction of classroom exchanges at the "micro level" consist of simple talk.

These examples conclude my discussion of the abstract structure of the theory as embodied in the structure of the model - what the pieces are, and how they fit together in general. In the next section, I start the move from the abstract to the concrete. The issue: If you are watching a teacher in action, how do you characterize what the teacher does in terms of the elements of the model? That brief discussion will be followed by a series of case studies that illustrate such characterizations.

4. FROM THE GENERAL TO THE SPECIFIC: BUILDING SPECIFIC MODELS OF INDIVIDUAL TEACHERS-IN-ACTION

We are now in a position to discuss the way in which the teaching of specific teachers (in particular contexts) can be modeled.

4.1. The basic idea

The basic idea is simple in theory. Say you are interested in modeling a teacher in a particular context - perhaps a student teacher working through a lesson on simplifying rational algebraic expressions (case 1, Section 5 of this paper), an experienced physics teacher teaching an innovative lesson of his own design (case 2, Section 6), or a mathematician kicking off his problem solving course (case 3, Section 7).

Suppose that you can identify or attribute:

* the teacher's beliefs relevant to this segment of instruction. These include the teacher's general pedagogical beliefs; beliefs about the discipline in general; about learning; about the specific content to be learned; and about the students.

* relevant goals regarding content, community, and individual students held by the teacher, at all levels of grain size. These range from very long-term goals (overarching pedagogical, content, and social goals for the whole of instruction) to very short-term goals (e.g., very specific content goals for specific subsegments of instruction).

* relevant aspects of the teacher's knowledge base. This means much more than identifying the knowledge the teacher has in the categories of content knowledge, general pedagogical knowledge, and pedagogical content knowledge (and perhaps others delineated above). It also means identifying the chunks in memory (routines the teacher might use, his or her content- related scripts, mini-lectures, etc.) in which that knowledge is potentially accessible for use by the teacher. These chunks are the components of which action plans are formed.

Suppose in addition that you can assign priorities and "conditions of activation and use" to all of the entities above. This assignment begins, of course, with the delineation of the teacher's lesson image. The lesson image specifies action plans, and one can elaborate the rationales for them and the beliefs that underlie them. But, there is more. For example, at what levels are various goals and beliefs typically activated in particular contexts? What kinds of events are likely to trigger changes in the levels of activation of those beliefs? Which action plans are typically associated with which constellations of goals? Which action plans typically get implemented under certain sets of constraints?

With this information, you are in a position to assemble the model of the particular teacher (in that specific context). You can now simulate the lesson by running the model. The initial state for the model has the activation levels of beliefs, goals, and knowledge, set to correspond to those in the teacher's lesson image. Then, you can run through the simulation of the lesson by having the model go through its action plans, updating the model with information about what has transpired. The model can be used to make predictions in the following way. At any point, you can posit a set of circumstances and ask how the model would respond. The model will work out a response as follows. If the circumstances you posit are accommodated comfortably by the current action plan, the model will propose the appropriate response according to that action plan. If the circumstances reflect something new - e.g., a student raises an issue not accommodated by the current action plan - then the reprioritization machinery described above clicks into motion. What beliefs does the event activate? Which ones are activated strongly enough to result in the modification of goal priorities? Which action plans might be selected that are in synch with the new constellation of strongly activated goals? Which of those are consistent with current lesson constraints? These considerations may result in the model's selection of a new action plan. If that happens, the prediction of the model is that the teacher would react similarly.

What is simple in theory may, of course, be rather complex in practice. One major issue, of course, is how one can identify, with some degree of confidence, the items listed above (beliefs, goals, and knowledge) and their activation levels in various classroom contexts. Part of the answer is to use as many sources of data as one possibly can, and to seek triangulation across them. Primary among the research group's data sources are tapes of the instruction itself. We analyze the tapes in fine-grained detail, looking for patterns and regularities. Such regularities may suggest that the teacher is employing specific routines. They may also, when a teacher consistently responds in certain ways to questions posed by the students (are questions encouraged? turned back to the students? answered briefly?), point to certain beliefs and goals that we may attribute to the teacher. We comb the tapes for overt statements of goals and beliefs, and try to match them to the teacher's actions. Ultimately we try, for each coherent chunk of classroom interaction, to identify the constellation of beliefs and goals to which that chunk of classroom interaction corresponds (see below for detail). The next major sources of information are the teachers themselves, and any artifacts they may be able to offer regarding instruction. Teachers may be interviewed before and after the instruction. They may participate in the ongoing analyses of the tapes (the amount of time devoted to the analysis of an hour of instruction is typically measured in months, sometimes years), and may be co-authors of the

analyses.[9] They may provide extensive documentation of before-and-after perceptions of instruction, by means of lesson plans, class logs, and papers about their actions and intentions. (For example, Jim Minstrell has written extensively about his teaching, in particular about the lesson discussed below; Deborah Ball has an extensive data base related to the tape we are analyzing as part of our ongoing work.) The goal of the research group is to use all of these sources to gather as coherent a picture as we can of the teacher's beliefs, goals, and knowledge.

The balance of this section provides a brief, generic description of how we analyze a lesson and of what we try to produce in the analysis. Having done so, I will turn to the case studies described in the following three sections.

4.2. Lesson parsing and model building.

This subsection of the paper offers a brief, abstract description of the process and representations we use for parsing and modeling a segment of instruction. One additional piece of terminology is necessary at this point. As discussed above, an action plan is a set of actions the teacher *intends* to take, in order to achieve a constellation of high priority goals. What the teacher intends and what actually takes place are often closely related of course, but they are not necessarily the same. In parsing a lesson we are interested in describing what actually took place. The segments of instruction we describe will be called *action sequences*. One major issue for analysis, of course, is how well the action sequences that take place correspond to the action plans held by the teacher - and what happens when they don't.

The parsing, which proceeds in stages, consists of the iterative decomposition of a body of instruction, which we shall refer to generically as a "chunk," into smaller chunks, each of which coheres on phenomenological grounds. In any analysis, the first chunk is the body of instruction being analyzed.

Figure 4, which represents an abstract parsing of a hypothetical lesson,[10] illustrates the form of the representation we use. On the very far left of Figure 4 one finds a complete line-by-line transcript of the instruction. The length of the instructional sequence is arbitrary. For example, Section 5 of this paper analyzes a lesson segment that lasted a few minutes; Section 6 analyzes a substantial part of a discussion that lasted a full class period; another analysis might take as its first chunk a unit that lasted a few weeks, or possibly a year's instruction.

---

[9]Note that Mark Nelson is a co-author of the paper that provides the substance of the analysis discussed in Section 5, and Jim Minstrell is a co-author of the paper that provides the substance of the analysis discussed in Section 6.

[10]This figure is adapted with permission from Schoenfeld, 1996.

Theory of Teaching in Context Draft EE Alan Schoenfeld 7/7/98

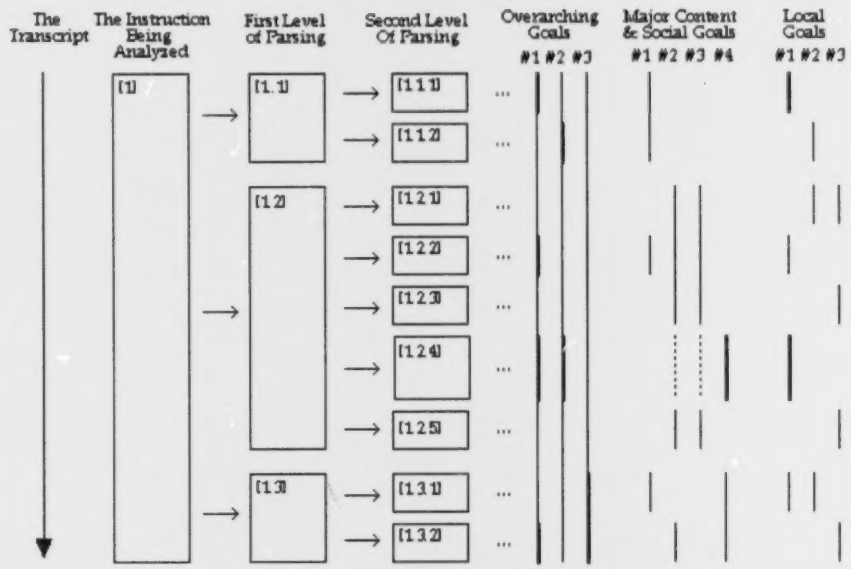

Figure 4
An abstract representation of the parsing of a body of instruction, including the transcript, the iterative decomposition of the instruction into chunks, and goal traces corresponding to those chunks

Starting in the second column of Figure 4 one finds the parsing of the instruction itself. The large box labeled with a 1, which is denoted [1], represents the instruction as a whole. To its right is the first level of parsing. To produce that parsing the first chunk is examined for break points. Typically a break point represents a change in the character of the instruction that is significant at the current level of grain size - that is, a change in focus, direction, emphasis, etc., that is notable with respect to the chunk of instruction being parsed. (Break points might correspond to the end of the discussion of a particular topic and the introduction of a new one, to the discussion of a problem, to a shift in classroom organization from whole-group to small-group, etc.) When the first level of analysis is completed, the result is a decomposition of chunk [1] into a sequence of sub-chunks, [1.1], [1.2], [1.3], etc. The parsing process is then repeated, with each of the chunks in the first level being further decomposed in a similar manner. The process continues down to a very fine level of detail, often resolving into line-by-line interactions between teacher and students.

In short, this aspect of the parsing produces a decomposition of what happened during instruction into a nested collection of action sequences. A small part of the underlying analysis that substantiates and justifies the decomposition is represented in the goal trace that appears on the right hand side of Figure 4.[11] As noted in Section 3 of this paper, the expectation is that there will be at least one strongly activated goal co-extensive with each chunk of instruction. Some overarching goals, such as the wish for students to perceive the subject matter as a sense-making activity, or for the classroom to function as a "community of inquiry," may be present at all times, but they may be highly activated during particular segments of instruction. This is represented in the first goal column in Figure 4, where, for example, Overarching Goal #2 is particularly strong during action sequences [1.1.2] and [1.2.4]. Other goals (say that students understand a particular "big idea") may recur throughout the instruction (cf. the column concerned with Major Content and Social Goals), and some (e.g., taking attendance) may be seen once at the beginning of the instruction and not again.

Of course, the goal trace represented in the right hand side of Figure 4 represents only the tip of the analytic iceberg. For each action sequence that is represented in Figure 4, we will want to know whether it was expected or unexpected, how (if at all) it corresponded to the teacher's lesson image, what "triggered" it, what beliefs might have shaped the way it took place, what goals the teacher's actions were intended to satisfy, what kinds of knowledge the teacher depended on in the interaction, and what brought the episode to a close. A formal component of the analytic scheme is that for each chunk that appears in the parsing of instruction, answers must be produced for each question on the list of questions given in Table 1.[12]

---

[11]The goal trace is read in the obvious way, as projections either rightward (from a chunk in the parsing to one or more one goal traces that are coextensive with the chunk, indicated that they are highly activated during the action sequence represented by that chunk) or leftward (where tracing backward from a goal to a chunk that is coextensive with it indicates that the goal is strongly activated during that chunk).

[12]This table is adapted with permission from Schoenfeld, 1996.

- Triggering event(s). What caused the action sequence to be instantiated?
- Beliefs related to the action sequence. (Beliefs, in context, shape goal prioritization and can serve to explain why a particular action sequence was selected or implemented.)
- One or more goals that the action sequence was intended to satisfy (with at least one such goal being at the grain size of the action sequence).
- Type of the action sequence. Was it a routine, a script, an improvisation, etc.?
- The kinds of knowledge that the action sequence taps into or depends on (e.g., specifics of subject matter knowledge or pedagogical content knowledge).
- The part of the lesson image to which the action sequence corresponds, if indeed it does.
- Terminating event(s). What caused the action sequence to be terminated?

Table 1
A delineation of the knowledge, goals, and beliefs to be specified for each of the action sequences in the parsing represented in Figure 4.

Given any sequence of instruction, much of the descriptive substance of the analysis is captured in the combination of the parsing represented in Figure 4 and the supplementary analyses contained in a full set of answers to each of the questions in Table 1 for each action sequence represented in the parsing. At minimum, this parsing-plus-detail provides a fine-grained characterization of what took place, and some good information regarding why - if the analysis is correct, the beliefs, goals, and knowledge attributed to the teacher should be internally consistent and should describe, down to a very fine level of detail, why the teacher chose to do particular things at particular times.[13] Moreover, this information provides the basis for prediction.

Let me constrain the notion of prediction carefully by example. It should be clear that in a general sense we cannot predict with high accuracy, on a fine-grained level, what people are likely to do in various situations - especially situations

[13]There are, of course, methodological issues to consider regarding how we develop a valid and reliable parsing procedure, and consistent rather than ad hoc answers to the questions in Table 1. As noted in Section 4.1, triangulation using multiple data sources is an important way to increase the likelihood of making valid interpretations. A research method known as competitive argumentation is one way we have of guarding against ad hoc explanation. (See Schoenfeld, Smith, & Arcavi, 1993, for a description of how we employ it for transcript analysis.)

that differ substantially from ones in which we have studied them carefully. For example, I have rather different goals and I utilize different parts of my knowledge base when I teach my undergraduate problem solving course and my research group, and I had different ones still when I served as a tutor for our work on modeling tutoring. To think that a deep understanding of my behavior in one of those contexts would allow you to predict my behavior in one of the others in any but the most vague and general way is silly, to say the least. On the other hand, I will teach my problem solving course again. As Arcavi, Kessel, Meira, and Smith (1998) can testify, there is great consistency in my goals and knowledge from one version of the course to another. It would be possible, I think, for Arcavi et al. to be given a tape of the first day of my new course, watch it up to a point where something interesting happens, and then predict how I would act next (in the sense of saying that on the basis of their analysis of my goals, beliefs, and knowledge, some actions are very likely, some are completely unlikely). Likewise, I could attempt the same for Jim Minstrell teaching his benchmark elaboration lesson, or for day (n + 1) of Deborah Ball's instruction, having carefully analyzed selected aspects of her course through day n .

The means for doing so is clear, and much of the necessary information is contained in the parsing-plus-elaboration described above. That is: In this context, what are the teacher's high priority beliefs, knowledge and goals? What has preceded this interaction, and what are the constraints on the teacher and class? At this moment, what is the teacher reacting to? Does this situation trigger the activation of any particular high priority beliefs? If so, what is the current constellation of beliefs likely to be in place? What goals are likely to be assigned high priority as a result? Given that set of goals and the teacher's sense of the students, which of the action plans that we know he or she has access to are most likely to be compatible with those goals and beliefs, and thus to be chosen?

This discussion concludes my general description of the model and related issues. In the next three sections of the paper I turn to four case studies, two described in detail and two summarized in brief. I begin with a discussion of a lesson segment taught by Mark Nelson.

5. CASE 1: MARK NELSON[14]

Here I provide the analysis of a short segment of instruction. My purpose is to illustrate the parsing scheme and to establish the basis for a comparative discussion (vis-à-vis the case discussed in Section 6) of the ways that lessons evolve as a function of the beliefs, goals, and knowledge that teachers possess.

5.1. Context

The lesson segment I analyze occurred in the middle of an Algebra 1 class taught by Mark Nelson. Nelson was a student teacher. He had been assisting and teaching this class, which takes place in an urban public high school, for several months. The course used a standard Algebra 1 text, and it followed the text closely.

The topic of the whole lesson was the arithmetic of exponents. The main content goal of the lesson was to enable the students to deal with the reduction of rational algebraic expressions such as $\frac{x^a y^b}{x^c y^d}$ where $a \geq c$ and $b \geq d$ into the form $x^{a-c} y^{b-d}$. A subsidiary goal was to deal with the special case $\frac{x^a}{x^a} = 1$.

Generally speaking, the class was on the noisy side but engaged – indeed, Nelson considered the noise to be a concomitant of engagement, and although he occasionally had to "shush" the class so that students could hear, the atmosphere was one of involvement with the content. Earlier in the lesson, the class had developed the basic idea of subtracting exponents when dividing powers of the same base. They had worked through examples such as the reduction of $\frac{4^5}{4^3}$ by expressing that fraction as $\frac{4 \cdot 4 \cdot 4 \cdot 4 \cdot 4}{4 \cdot 4 \cdot 4}$ and canceling, yielding 4^2. Further work with examples such as $\frac{x^6}{x^3} = x^3$ led to the generalization $\frac{a^m}{a^n} = a^{m-n}$, when $m > n$.

Having worked through these examples with the whole class, Nelson had the students work in small groups for a while. They were assigned the problems

(a) $\frac{r^n}{r^n}$ (b) $\frac{x^a y^a}{x^b y^b}$ (c) $\frac{x^5}{x^5}$

and told that the class would discuss them when they had finished working on them in groups. As the students worked, Nelson circulated through the room

[14]The material presented in this section is based on the analysis given in Zimmerlin & Nelson, 1996.

observing their work. He then convened the class to discuss the problems. My detailed analysis starts at this point in the lesson.

5.2. Nelson's beliefs, goals, lesson image, and knowledge base.

Beliefs

Describing Nelson's beliefs - most importantly, his stance toward the discipline of mathematics and its teaching - is a delicate issue. To establish a context for the claims to follow, here is a discussion of a comparative analysis of videotapes of American and Japanese teachers.

> In our tapes, we have observed two strategies that American teachers employ to deal with incorrect responses to questions. One of these is simply to ignore the response and to let another student answer. This is consistent with the idea . . . that the American teacher's goal in asking questions is to get the answer she needs to move the lesson along. The same purpose is accomplished as well by the second strategy, that of reinterpreting incorrect or inappropriate answers so that they fit the teacher's expectations. . . . Implicit in the way American teachers ask questions is that the teacher knows the answer and is the authority on what is right or wrong. (Stigler, Fernandez, & Yoshida, 1996, pp. 240-241)

It can be argued that the belief system described by Stigler et al. plays a major role in shaping not only a teacher's responses to student questions, but the very nature of a teacher's knowledge base. If a teacher sees the primary use of student comments as that of providing springboards for his or her explanations, then there is little perceived need to delve deeply into the nature of the students' understandings as represented by those comments. Hence (in contrast to Japanese teachers - see Section 5.4) the teacher may not develop a diagnostic repertoire or skills at elucidating student conceptions of the subject matter. As we shall see, Nelson fits the pattern of Stigler et al.'s observations.

Goals and lesson image

Nelson's formal lesson plan was very sketchy (Zimmerlin & Nelson, 1996), but extensive interviews and conversations revealed that his lesson image and his goals for the lesson were in fact quite detailed. At the level of overarching goals, Nelson had content goals (to help students see where algebraic notations and procedures come from, and why algebraic rules are true) and social goals (to develop a classroom atmosphere in which students contribute to classroom activities). One sees those goals play out respectively in the inductive approach Nelson took to the content (the class worked through examples before codifying the results of their work as rules) and in the engaged character of the classroom interactions. More local goals for the discussion of the individual problems are given in the discussion of Nelson's lesson image, which follows.

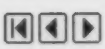

Nelson expected problem (a) to cause the students no difficulty, for it was essentially identical to the work the class had done in the previous part of the lesson. (In addition, by the time he convened the class to discuss the problem, he would also know that the small groups had obtained the right answer, confirming his judgment that this was an easy problem.) He expected to ask the students to call out their answers to the problem, to confirm the correct answer, and to check for problems. None were expected, and he would then move on to problem (b). (Specific goals for problem (a): to build students' confidence by having them arrive at the correct answer; to reinforce the procedure and notation for reducing rational expressions that have a single base.)

Problem (b) represented a slight extension of what the class had done before, there being two bases (x and y) rather than just one in the expression that he asked the students to reduce. Nelson thought that the students would have little difficulty with the problem. As in (a), he expected the students to call out their answers. In addition, he planned to ask the students for explanations about how they had arrived at their answers, with the intention of codifying the explanations and articulating them as a procedure to be used in general. (The specific goal for problem (b): to establish a variant of a known technique, expanding known results from single to multiple bases.)

In contrast, Nelson expected some initial student confusion with regard to problem (c). He planned to begin as with problems (a) and (b), soliciting student answers. There would, he thought, be some different answers, and some students who said they didn't know or understand. To deal with the confusion Nelson planned to work through the example at the board, expanding $\frac{x^5}{x^5}$ as $\frac{x \cdot x \cdot x \cdot x \cdot x}{x \cdot x \cdot x \cdot x \cdot x}$, canceling the xs, and obtaining $x^0 = 1$ as a result. This done, he would proceed to discuss the special case where $x = 0$, noting that division by 0 is not permitted so that 0^0 is undefined. (Specific goal for this problem: to build an understanding of zero exponents. Subgoals: to deal with the case where the base is non-zero, and then where it is zero.)

Knowledge base

Here is a brief description of the knowledge structures (subject matter knowledge, general pedagogical knowledge, and pedagogical content knowledge) on which it depended. Content-wise, this was familiar territory for Nelson - the mathematics is elementary for someone with his background. Pedagogically speaking, Nelson expected to use a questioning strategy in which he would ask students for the answers to the problems he had given them and then provide feedback on their responses, perhaps elaborating as well. This kind of questioning sequence, called an IRE sequence - the teacher *I*nitiates the

sequence with a question, a student *R* esponds, and then the teacher *E* valuates the student's answer - is quite common, as documented by a wide range of studies (see, e.g., Cazden, 1988; Lemke, 1990; Mehan, 1979). This approach provided a structure within which Nelson could expect to keep the students engaged, and in which the students would provide him with the fodder for his summary comments. Finally, given that this was his first time teaching this particular material, Nelson could not be expected to have a reservoir of pedagogical content knowledge for this topic (e.g., the specific awareness of potential trouble spots for students and how to deal with them). Such knowledge is developed with experience, if it is developed at all.

5.3. What took place

The transcript of the lesson segment, which lasted about four minutes, is given in Appendix A. Here I focus on main themes.

The discussion of problem (a) took just six lines of exchange and went as planned. It consists essentially of two IRE sequences, first obtaining the right answer and then confirming that there were no problems. (Note: The exchange in lines 5 and 6 is casual and friendly, an indication of the relationship Nelson has built with the students.)

The discussion of problem (b) proceeded similarly. The IRE sequence in lines 7-9 produced the answer xy to the problem. Nelson then asked "Why did you get xy?" The difficult-to-hear response included the word "subtract," to which he responded "OK." He then initiated the next sequence: "What did you subtract?" The student response "3 minus 2" provided the grounds for a quick "OK" and an elaboration: "you looked at the xs and subtract 3 minus 2. That gave you x to the first [writes x on the board] . . . and you looked at the ys and said 7 minus 6, that gave you y to the first [writes y so that $\frac{x^3y^7}{x^2y^6} = xy$ is on the board]. To this point, everything proceeded as planned.

Things did not proceed as smoothly with problem (c), in which Nelson ran into some unexpected roadblocks. The initial set of responses to his opening question "OK, (c). What did you get here?" was varied, as he expected. This provided the grounds for his working through the expansion of $\frac{x^5}{x^2}$ as $\frac{x \cdot x \cdot x \cdot x \cdot x}{x \cdot x \cdot x \cdot x \cdot x}$, and the subsequent canceling of xs by means of making slash marks through the numerator and denominator. His expectation was that the next question sequence, which began with "What am I left with?" (line 38), would produce the right answer – "1" – and that he could then summarize and wrap things up. But, the students did not respond as he expected. They called out answers that included "x," "zero," "zero over zero," and "nada."

Nelson, no longer on familiar ground, had to deal with unexpected confusion on the part of the students. To do so he generated an alternative impromptu example that involved canceling numbers instead of symbols.[15] With the goal of clarifying that any non-zero quantity divided by itself gives one - the basis for the solution to problem (c) - he said (lines 42 ff.), "Just a second. Just a second. What if we had 5 over 5? What's that equal?" The students did indeed say "1," but they failed to make the connection that he wanted them to make - that the canceling applied to the x 's as well.

For the next few minutes Nelson tried, without success, to rectify the situation. In line 45 he pointed out that canceling the 5's gave a 1 (and not "nada" as the students had suggested for the canceled x's). He tried again in line 48, pointing out that the answer is 1 in the case of $\frac{5}{5}$, and asking (expecting an affirmative answer) if the answer would also be "1" in the case of $\frac{x \cdot x \cdot x \cdot x \cdot x}{x \cdot x \cdot x \cdot x \cdot x}$. The students did not see the parallel. In the exchange that followed, none of the students produced an answer that Nelson could use as a springboard for an explanation – in lines 49 through 73 he and the students essentially talked past each other. Nelson was stymied. He had hoped to elicit the correct answer from a student and then to elaborate the method by which the answer was obtained. Without the right answer, he was stuck. In line 75, as a means of getting back on track, he called on a student whom he believed had obtained the right answer to the problem. The student said that the answer is "x to the zero power over 1." Interestingly, Nelson mis-heard this[16] as "x to the zero power equals 1," which is a succinct statement of what he needed to hear. This "correct" answer provided the means by which he extricated himself from his difficulties. In lines 76 though 87, he asserted that "5 over 5 equals 1, so this $-\frac{x \cdot x \cdot x \cdot x \cdot x}{x \cdot x \cdot x \cdot x \cdot x}-$ is going to equal 1 as well" and reminded students that when you subtract exponents (in the original expression $\frac{x^5}{x^5}$) you get x^0. He concluded by writing

$$x^0 = 1$$

on the board and telling the students "Get this in your notes . . . any number to the zero power equals 1."

---

[15]This is a very reasonable ad hoc pedagogical move. Since the students were having trouble canceling symbols, he drew them back to familiar ground - canceling numbers. However, his choice of the number "5" was unfortunate, in that those students who did not understand his rationale for returning to the domain of numbers might think that he was now doing something new with the *exponent* 5 from the original problem.

[16]Nelson told us that he heard the student saying "x to the zero power equals 1," and he was surprised to hear what the student actually said when he listened to the videotape of the lesson.

5.4. Analysis

My purpose in discussing this lesson segment is two-fold. First, I shall offer a diagrammatic parsing of the segment, to illustrate the representation we use and some of the issues it addresses. Second, I discuss the beliefs, goals, and knowledge that Nelson brought to the lesson segment, and how interactions among them affected the way the lesson segment played out.

Figure 5 provides the first few levels of a parsing of the instructional sequence, along with a parallel goal analysis (ellipses indicate where further decomposition could be undertaken). A complete analysis of this segment would call for extending Figure 5 to the level of line-by-line interactions, and also filling out a version of Table 1 for each of the boxes in Figure 5. As suggested by the discussion in Section 5.2, it would be a relatively routine matter to supply most of that detail. In what follows I forego that level of analysis, but point instead to some of the most interesting issues related to the lesson segment.

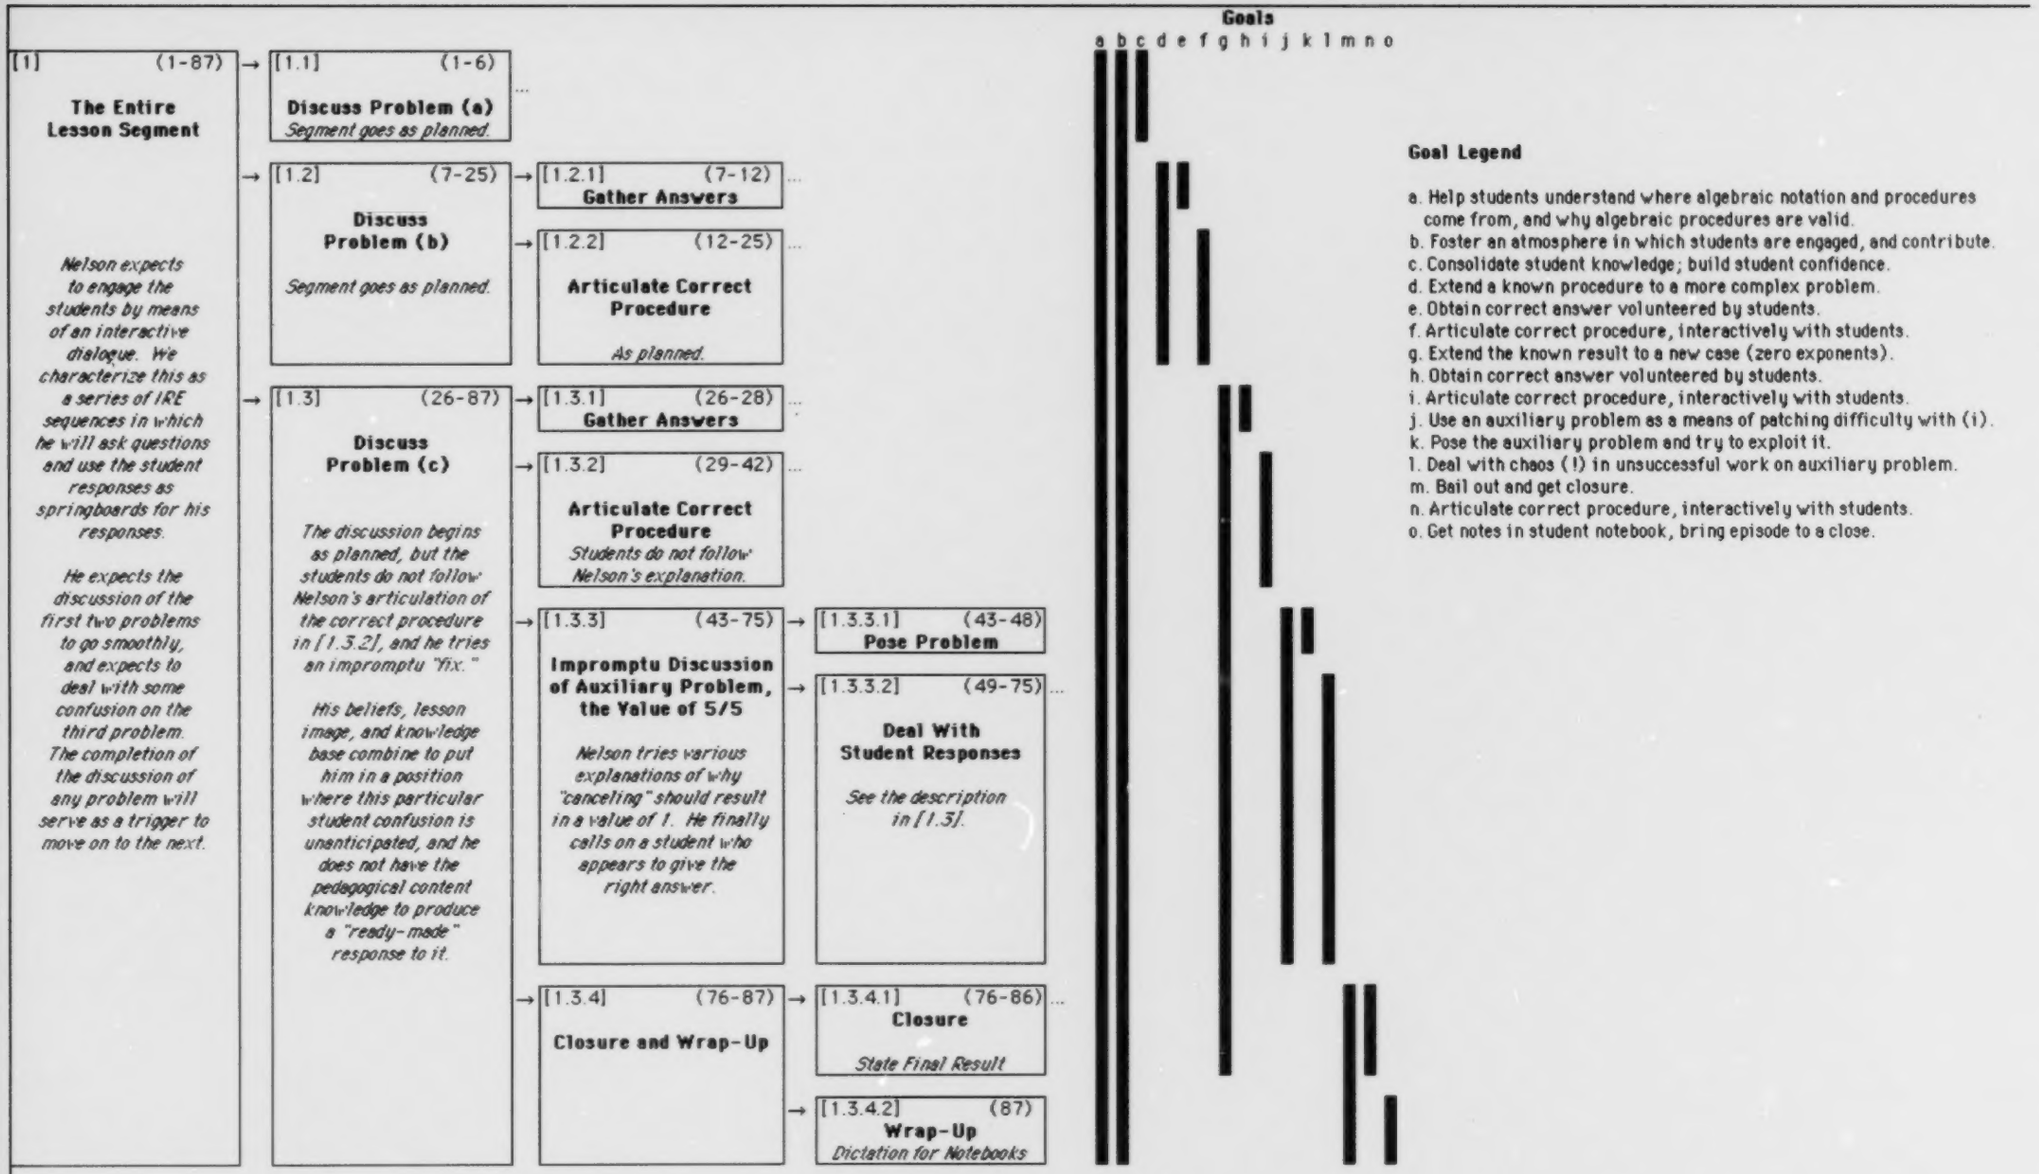
[1] (1-87)
The Entire Lesson Segment
Nelson expects to engage the students by means of an interactive dialogue. We characterize this as a series of IRE sequences in which he will ask questions and use the student responses as springboards for his responses.
He expects the discussion of the first two problems to go smoothly, and expects to deal with some confusion on the third problem. The completion of the discussion of any problem will serve as a trigger to move on to the next.
[1.1] (1-6)
Discuss Problem (a)
Segment goes as planned.
[1.2] (7-25)
Discuss Problem (b)
Segment goes as planned.
[1.2.1] (7-12)
Gather Answers
[1.2.2] (12-25)
Articulate Correct Procedure
As planned.
[1.3] (26-87)
Discuss Problem (c)
The discussion begins as planned, but the students do not follow Nelson's articulation of the correct procedure in [1.3.2], and he tries an impromptu "fix."
His beliefs, lesson image, and knowledge base combine to put him in a position where this particular student confusion is unanticipated, and he does not have the pedagogical content knowledge to produce a "ready-made" response to it.
[1.3.1] (26-28)
Gather Answers
[1.3.2] (29-42)
Articulate Correct Procedure
Students do not follow Nelson's explanation.
[1.3.3] (43-75)
Impromptu Discussion of Auxiliary Problem, the Value of 5/5
Nelson tries various explanations of why "canceling" should result in a value of 1. He finally calls on a student who appears to give the right answer.
[1.3.3.1] (43-48)
Pose Problem
[1.3.3.2] (49-75)
Deal With Student Responses
See the description in [1.3].
[1.3.4] (76-87)
Closure and Wrap-Up
[1.3.4.1] (76-86)
Closure
State Final Result
[1.3.4.2] (87)
Wrap-Up
Dictation for Notebooks
Goals
a b c d e f g h i j k l m n o
Goal Legend
a. Help students understand where algebraic notation and procedures come from, and why algebraic procedures are valid.
b. Foster an atmosphere in which students are engaged, and contribute.
c. Consolidate student knowledge; build student confidence.
d. Extend a known procedure to a more complex problem.
e. Obtain correct answer volunteered by students.
f. Articulate correct procedure, interactively with students.
g. Extend the known result to a new case (zero exponents).
h. Obtain correct answer volunteered by students.
i. Articulate correct procedure, interactively with students.
j. Use an auxiliary problem as a means of patching difficulty with (i).
k. Pose the auxiliary problem and try to exploit it.
l. Deal with chaos (!) in unsuccessful work on auxiliary problem.
m. Bail out and get closure.
n. Articulate correct procedure, interactively with students.
o. Get notes in student notebook, bring episode to a close.

Figure 5
Nelson's Discussion of Problems a, b, and c.

Theory of Teaching in Context Draft EE Alan Schoenfeld 7/7/98

Issues of beliefs and epistemology

I want to stress that Nelson wanted his students to understand the material and that he structured the lesson accordingly. His lesson structure was inductive, in that he expected students to work a series of examples before he would codify the methods that solved them as rules. And, Nelson planned to use the canceling in problem (c) as a vehicle for explaining why $x^0 = 1$. So he was certainly not demonstrating procedures and asking students to follow them in rote fashion.

Nonetheless Nelson's lesson, in interesting ways, does not delve very deeply into student understanding. Student comments are an important component of the lesson, but they are important in that their primary role is to serve as contributions to, or stimuli for, the teacher's explanations. Divergent ways of thinking are not conceived of as windows into ways of thinking about the problem domain. Indeed, they are not anticipated. Nelson is not alone in this, either as a new teacher or as a teacher in the United States. As the quotation from Stigler, Fernandez, and Yoshida (1996) in Section 5.2 indicates, a focus on "what I will do in the classroom" is typical of the planning of the U. S. teachers they studied. And, what the teachers will do is explain why something works. This approach stands in rather stark contrast with the lesson planning focus of the Japanese teachers studied by Stigler et al. Those teachers tended to frame their lesson plans in terms of expected student understandings. Moreover, they conceptualized their own activities in the light of how they intended to shape or respond to the kinds of understandings student comments revealed. (See also Ma, 1996, for some parallels with regard to teachers in China.) In fact, the support materials for Japanese curricula often contain descriptions of typical (correct and incorrect) student responses to problem situations. Had Nelson been so prepared - both in the epistemological sense and in the more narrow sense of being aware of typical student responses - he might not have been caught off guard as he was.

Issues of knowledge - a more detailed analysis

Relatively new to teaching, Nelson had not yet developed much of the pedagogical content knowledge that is relevant to this topic – knowledge of typical student understandings and confusions in this area, and a repertoire of "fail-safe" examples and techniques for dealing with student difficulties. As a result, he ran out of pedagogical options rather fast. Consider chunk [1.3] of the lesson. To recapitulate in brief, chunk [1.3.1] went according to plan, and the transition to chunk [1.3.2] took place as expected. Nelson then began the decomposition of $\frac{x^5}{x^5}$. He expanded the numerator and denominator to obtain the expression

(*) $$\frac{x \cdot x \cdot x \cdot x \cdot x}{x \cdot x \cdot x \cdot x \cdot x}$$

and, when the students told him to cancel, he did so. Students and teacher then faced the expression

(**) 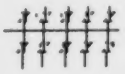

on the board. To Nelson, this clearly represented the expression

multiplied by itself five times, and was thus a product of 1's. This is what he expected the students to see, and he was caught short when they did not.

Why didn't they, and would it have been possible to anticipate the fact that they wouldn't? It is possible, though by no means guaranteed, that a teacher with more experience and the predilection to explore student understandings would have recognized that the reasoning summarized in the previous paragraph could represent a significant leap for the students. Note than when Nelson canceled x 's in the expression above, the pairing of each x in the numerator with a corresponding x in the denominator, and the creation of a "1" by slashing through the fraction, was purely implicit. When he performed the canceling operation, making a series of slash marks through the expression (*), the students may have seen just that - the teacher making a series of slash marks through the expression. Imagine for example that you are algebraically naive and are looking at the "canceled" expression (**) above. It looks like everything is simply crossed out. Hence when you are asked what's left, "zip" or "nada" is a perfectly reasonable response.

This is potentially a very serious problem, for students who fail to see the factorization represented by the canceling operations would find it nearly impossible to make sense of what followed. In addition, this was only one of two parts of the argument Nelson was implicitly relying upon. The second part, that each of the canceled terms has a value of 1, was also problematic. Nelson clearly intended that the students to make that connection – that is why he asked the question "What if we had 5 over 5? What's that equal?" (line 43) in response to the students' confusion about the value of the expression (**). However (see lines 43-52) the fact that he was trying to draw a parallel between the term $\frac{5}{5}$ and the term $\frac{x}{x}$ was never clearly articulated. In consequence, the students might not have realized that he was referring to terms of the form $\frac{x}{x}$ when he pointed to (**) and asked, "Is there a 1 there?" (line 48).

Given that much of Nelson's argument was never made explicit to the students, it is not surprising that the students would find it difficult to follow. A teacher with more experience might have recognized the potential for such difficulties, and taken steps to deal with them.

No criticism of Nelson is intended here. Nor is there the suggestion here that any beginning teacher should be able to unpack the complex chain of reasoning described above - especially on the fly, in response to unexpected student confusion. It is only plausible to expect that teachers could do so if they had many of the pieces (among them a predisposition to use students' comments as windows into their understanding, knowledge of what students are likely to understand, knowledge of how the task can be taken apart, of which connections might need to be made explicit, and of which pedagogical approaches are likely to be productive) in place. A beginning teacher is most unlikely to have access to all this, which comes with experience and reflection if it comes at all.

To sum up in brief, Nelson's epistemological stance (which did not predispose him to anticipate divergent student responses) and his knowledge base (solid on content but lacking in awareness of student perceptions and how to respond to them) combined to put him in the bind reflected in lines 49-75 of the transcript. When the students failed to generate what he needed, he was stuck. His on-the- fly attempts to clear things up didn't work. Then, in line 74 he called on a student whose paper had the right answer, "1". Mis-hearing what the student said in line 75 as "$x^0 = 1$" provided a way out, which he took. (Note that Nelson's actions here were consistent with the strategy described by Stigler et al., "reinterpreting incorrect or inappropriate answers so that they fit the teacher's expectations.") He brought this segment to a close with a standard move: he told the students what they needed to have in their notebooks.

5.5. Could we model this lesson segment in detail?

I would like to argue that much of the substance needed for creating a very detailed model of the instructional segment given in Appendix A is already present in this discussion. As discussed in Sections 5.2 and 5.4, we have a good sense of Nelson's beliefs, which focus on developing understanding through teacher explanation rather than on the elucidation of student conceptions. Similarly, we have a good sense of Nelson's goals and lesson image; we know and can model the mechanism he will use for interacting with the students, IRE sequences.[17] We know as well that he expects student confusion when he asks for their answers to question (c), and that he expects to lead the students through a solution. Further, we know that his store of pedagogical content knowledge

[17]We do not claim that Nelson is consciously using IRE sequences, only that such question-and-answer sequences provide a very consistent and accurate description of the way he interacts with students. Hence IRE sequences can be used both to describe and predict Nelson's behavior in class.

related to this topic is on the weak side, and that his last-ditch strategy for dealing with situations is "telling," preferably in elaboration or response to the correct answer provided by a student.

With that information, almost all of Nelson's actions can be explained. Note (as discussed in Section 5.2) that the first part of the lesson unfolds almost exactly as the model predicts.[18] Nelson uses IRE sequences to deal with problem (a), then makes the transition to problem (b) as planned and deals with that problem in the same way. He makes the transition to problem (c) as the model suggests, and proceeds precisely according to plan until, in line 39, the students do not provide the expected response, "1", to the question he has posed. At this point the model predicts that (a) because of his belief structure, Nelson would not have ready access to alternative strategies (such as exploring the students' understandings), and (b) because of his knowledge base, Nelson would not have ready access to the pedagogical content knowledge that would enable him to deal comfortably with this situation. In consequence, he would have to deal with this situation in an ad hoc fashion. The model could not predict how he would deal with it. (One could provide a post hoc rationalization of his choice of 5/5 as an example, but that is a different matter.) However, knowing that Nelson has taken the direction he has, the model would predict that he would soon find himself in a bind, since he would not understand the sources of the chorus of incorrect responses that came from the students. Finally, when Nelson is truly stuck (lines 75 ff.), he employs the exit strategy that the model suggests: once again using an IRE sequence, he calls on a student who can provide the answer he needs to hear and then uses the student's response as a springboard for a declarative presentation that brings the session to a close.[19]

I conclude this section with some brief comments regarding the main components of the model and their fit with the data analyzed here. Note that Nelson's beliefs (his epistemological stance, which focuses on the central role of teacher explanation rather than on the importance of elucidating student conceptions), his goals (both content and social), his lesson image (which involves soliciting answers from students and using them as springboards for elaboration) and his knowledge base (which provides him access to general content knowledge related to the lesson, general pedagogical strategies such as the use of IRE sequences, but little specific pedagogical content knowledge for this subject area) all fit together as the model predicts. At the beginning of the

---

[18]For reasons of space constraints I provide a brief summary of what happened rather than a line-by-line analysis of the Nelson transcript, accompanied by a full description of goals and action plans, etc. The analysis at that level is rather straightforward and I invite readers to work through the transcript at that level of detail.

[19]Of course, there is no way to predict that Nelson would mis-hear the student, and we cannot know what would have happened if he had heard the student correctly. (The absolute last-ditch strategy is to bring things to a halt and simply announce the correct answer, and there is some chance he might have done that - but that is mere conjecture, and we have no way to substantiate it.)

lesson segment his high priority beliefs, content goals (specifics related to dealing with exponents), and social goals (making sure students are involved) are all in synch, and entirely consistent with his lesson image. Nelson's knowledge base provides the resources for proceeding as planned. His questioning strategy provides an opportunity for student engagement, for getting answers "out in the open," and for providing explanations at appropriate points. Indeed, it is the tight synchrony of beliefs, goals, and lesson image that becomes problematic later on. In combination, they have led Nelson to focus on one particular pedagogical approach. When the questioning strategy proves unsuccessful and his knowledge base does not provide access to alternative strategies for responding to student confusion, Nelson is then compelled to deal with this difficulty on an ad hoc basis. When it is clear that the ad hoc methods are not successful, the (emergent) goal of bringing the lesson segment to a close is assigned a very high priority. Nelson chooses a path of action that is consistent with this goal and his other high activation beliefs and goals as well. He calls on a student whom he believes will provide the right answer, and responds to what the student says by producing the explanation (drawn from his knowledge base) that he wants the students to get into their notes.

In sum, this analysis does a good job of explaining what Nelson did and why, at a level of mechanism, in terms of the main constructs in the model. Indeed, were we to perform the gedanken-experiment of stopping the tape and asking what he would do next, it would do a good job of predicting his actions a significant proportion of the time.

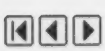

6. CASE 2: JIM MINSTRELL[20]

The case examined in Section 5 focused on a beginning teacher conducting a traditional lesson grounded in a standard text. This second case moves to a very different arena, focusing on a highly experienced teacher conducting an innovative lesson of his own design. This move to a radically different example is part of a deliberate research strategy, a means of seeing whether the model can accommodate teaching that varies widely across multiple dimensions - various teachers' beliefs, lesson images, goals, and knowledge bases. Should the model work in this case as well (that is, should we be able to capture these two very different teachers with the same basic architecture of the model), then there would be reason to think that the structure of the model could indeed be general and widely applicable. On the other hand, running into serious obstacles would point to places where refinement of the model is necessary - or, in the worst case, major difficulties would reveal that the basic architecture of the model is flawed and we would be forced to abandon it.

The analysis that follows focuses on a segment of a lesson that has some surface parallels to then Nelson lesson segment. Each of the lessons was proceeding according to plan when a classroom event resulted in the teacher making an extended impromptu excursion into unplanned territory. In each case the move into new territory was grounded in, and consistent with, the teacher's beliefs. And, in each case the way in which the discussion evolved was very much a function of the teacher's knowledge base. Hence this lesson, in addition to providing a new test case for the model, also provides data for an interesting compare-and-contrast discussion with the lesson segment described in Section 5.

6.1. Context

Jim Minstrell, who teaches physics at Mercer Island High School in Washington State, has been the recipient of the Presidential Award For Excellence In Science Teaching. He has pioneered the development of high school physics instruction that focuses on sense-making, and has written and lectured extensively on the topic (see, e.g., Minstrell, 1982, 1989, 1992;
Minstrell & Stimpson, 1996; van Zee & Minstrell, 1997a). For a discussion of the lesson analyzed here, see van Zee and Minstrell (1997b).

Minstrell teaches a series of lessons of his own design. These are intended to engage students with physics in a way that is meaningful to them and that conveys the spirit of the discipline as he understands it. I shall analyze part of a lesson that took place at the beginning of the school year. In this lesson the goals of engagement and sense-making (as well as the goal of creating a community of

[20]The material presented in this section is based on the analysis given in Schoenfeld, Minstrell, & van Zee, 1996.

thinkers, where students take an active role in figuring things out) play out in very strong ways.

The instruction I discuss took place the fourth day of the year. The first two days of the course had been devoted to preliminaries - playing the "name game" on Day 1, which culminated at the end of every class with Minstrell calling on every student by name, and the administration of a diagnostic test on Day 2 to document the students' initial knowledge about physics. Instruction began on the third day, when Minstrell introduced the Blood Alcohol Content problem, a problem he had designed "to engage teenagers in thinking about measurement issues in a meaningful everyday context." The gist of the problem is this. The blood alcohol level of a person arrested for drunk driving had been measured by five individuals who had obtained five different numbers. How could one deal with those data in a sensible way, and arrive at a "best value" for the driver's blood alcohol level?

Discussions of the Blood Alcohol Content problem occupied most of the third day. One focal point of the classroom conversation was: Which numbers should count? What if one of the five measurements taken was discrepant with the others? What if one was taken by a doctor or some other "expert"? Should the first be discarded, should the second count more in some sense? The class explored reasons for considering some data and not others (e.g., high and low scores are thrown out in the judging of ice skating events in the Olympics) and the idea that choosing such numbers is a matter of discretion and judgment, not just following rules. The class did some of its work in small groups. As preparation for the next day's lesson, one representative of each group independently measured the length and width of the same table. The students were given the homework assignment of (a) finding the best value for the blood alcohol content and deciding whether or not the driver was drunk, and (b) calculating best values and uncertainties for the length and width of the table. This background provided the context for the fourth lesson, which Minstrell refers to as a benchmark elaboration lesson. The focal point of the conversation was: Once you have decided which measurement values to keep, how should you combine them to obtain a "best value"? In pursuit of this issue, the class explored three different ways of combining the data - the three standard measures of central tendency, mean, median, and mode. Students discussed the fact that each of these procedures may give a different number, which may be more or less meaningful given the situation at hand.

As discussed in detail in Section 6.3, much of that lesson went according to plan. The review of which numbers they might choose, and why, covered the desired territory. The first part of the discussion of different ways to combine numbers, which covered mean and mode, also took place as expected. Then a student suggested a novel method for computing the average value. This suggestion led Minstrell and the class into unexpected territory. They explored the student's proposed method and how it was related to the standard definition of average

for about eight minutes. This discussion completed, Minstrell returned to his lesson plan. He wrapped up the consideration of "best value." Under some time pressure at that point because of the unplanned eight-minute discussion, Minstrell then led the class through a discussion of measurement error and the reporting of significant digits. The lesson closed with the assignment of homework for the next day.

My analysis will focus most closely on the discussion of best value, in particular on Minstrell's unplanned exploration of the student's suggestion. An extensive analysis of the whole lesson may be found in Schoenfeld, Minstrell, and van Zee (1996); see also van Zee and Minstrell (1997).

6.2. Minstrell's beliefs, goals, lesson image, and knowledge base

Beliefs and goals

Minstrell conceives of physics as a sense-making activity, and he believes that students' classroom experience with physics should support the development of students' abilities as sense-makers. He encourages open and free discussion in class and has the class spend a great deal of time thinking things through. In order for students to do the kind of sense-making he would like, they need to engage with the material and not simply be told what works. Hence Minstrell works to create a discourse community that involves small-group work among students, and he structures teacher-student exchanges in a way that minimizes "telling." As part of his commitment to student inquiry, Minstrell values and places very high priority on pursuing substantive ideas raised by students - indeed, pursuing important ideas raised by students will tend to get higher priority from Minstrell than following a predetermined lesson plan. Minstrell also believes in the value of reflectiveness and works to promote it. Questions regarding why something makes sense, or how it fits with other things the class knows, will be encouraged and pursued.

At a very general level, then, Minstrell's goal is to foster the development of a sense-making community. His specific content-related goals for this lesson concern the blood alcohol content problem and the measurement of the table. Minstrell wants the students to emerge from instruction with a sense of the issues related to which numbers count, how to combine the numbers, and how to think about measurement error and precision in reporting. The students should also have knowledge of specifics such as the three measures of central tendency.

Lesson image

Minstrell has taught the Blood Alcohol Content lesson many times. He has a very rich and detailed lesson image, one that is deeply coupled with this

pedagogical style. In rough outline, Minstrell expected the lesson to have three main segments:

(1) a relatively short introductory segment intended to deal with bureaucratic necessities and any questions the students might raise,

(2) a discussion of "best number" (deciding which data to use and deciding how to combine them) that would use as a case in point the measurements of the table's length and width that the students had taken the previous day,

(3) a discussion of "precision" (significant digits and measurement error) that would be grounded in a discussion of the value that the students reported for the best number in part (2),

followed by a brief wrap-up and a homework assignment.

Minstrell has a straightforward script for dealing with the first segment. He will say what needs to be said, ask for questions, respond to them, ask again, and proceed only when it is clear that students have run out of questions.

For the second and third segments Minstrell expects to employ a questioning method that he calls the "reflective toss." (Minstrell speaks of "'catching' the meaning of the student's prior utterance and 'throwing' responsibility for thinking back to the students, not only to the individual student doing the talking but also to all of the students in the class" (van Zee & Minstrell, 1997, p. 227).) The basic idea is to have the students generate, as much as possible, the intellectual content of the conversation. Minstrell will start the conversations with a question ("What methods do we have for dealing with this?") and follow it up with a request for more detail ("How does that work? Why does that work?"). When students make statements that he thinks need clarification or elaboration, Minstrell will typically ask the student or the class to comment on what has been said, or to provide more detail. However, he stands ready to provide it himself if the students do not provide what he needs.

A detailed analysis suggests that Minstrell's use of reflective tosses is a component of a larger pedagogical approach called *interactive elicitation* (Schoenfeld, Minstrell, & van Zee, 1996), which works as follows. In the review of which numbers should count and why, for example, Minstrell will ask the students to remind him of the reasons that they had discussed for considering some rather than all of the measurements. When a reason is raised, he will ask for clarifications and additional detail, as described in the previous paragraph. When he feels that the approach under discussion has been dealt with adequately, he will then say "Are there any others?" and repeat the clarification process. This sequence will be repeated until the well runs dry. If the students stop short without generating the list of reasons he expects, he will suggest that

there are others; if the students are unable to generate any others he will mention one (and then turn to them for clarification). Indeed, Minstrell will ask "Any more?" even when the previously generated list has been re-generated, and he will wait a very long time (more than 10 seconds, which in most U. S. classrooms seems an eternity) before saying "Yes, I think that's the list we had." He will conduct the discussion of how one might combine the selected numbers similarly. I note that interactive elicitation is used by Minstrell not only for review sessions but for discussions of new material as well. For example, the discussion of precision that Minstrell envisions for the latter part of the lesson will follow the same general pattern - but because some of this material is new his questions will be more focused and he will make more substantive contributions to the conversation.

What the preceding discussion does not convey is the richness of Minstrell's lesson image and his comfort with what is to happen. Having taught this material many times before, Minstrell can anticipate most of what the students will suggest and how he will react to what they will say. For example, if a student makes a comment that calls for clarification or elaboration, the odds are that Minstrell will have heard many similar comments before and that he can (re-?)generate an appropriate reflective toss without much effort. The richness of Minstrell's lesson image and his ability to implement it rest on his knowledge base, which I now describe briefly.

Knowledge base

I shall briefly discuss three categories of relevant knowledge: Minstrell's subject matter knowledge, his general pedagogical knowledge, and his pedagogical content knowledge. I then discuss how they interact.

Minstrell knows this territory very well and has quick (mental) access to a wide range of content that enables him to deal easily with issues that arise in conversation with the students. Here is one example. I mentioned above that in the course of the lesson a student made an unexpected suggestion for obtaining the "best number." The student said:

> "This is a little complicated but I mean it might work. If you see that [a specific value for the width of the table] shows up four times, you give it a coefficient of 4, and then [a different value] only shows up one time, you give it a coefficient of 1, you add all those up and then you divide by the number of coefficients you have."

One of two possible interpretations of this statement (dividing a weighted sum by the sum of the weights) is, of course, a standard way to represent the average of a collection of numbers. It may seem a trivial point, but the fact that Minstrell can recognize it immediately as such, and then use his knowledge of this and related things, makes it possible for him to proceed smoothly with the lesson. If

he had needed to stop to figure out what the student said and what the mathematical entailments of that suggestion were, the lesson might have taken a very different turn. [N. B. The transcript of this part of the lesson is given in Appendix B. It is discussed extensively in Sections 6.3 and 6.4.]

A second aspect of Minstrell's knowledge base is his general pedagogical knowledge. Minstrell is a highly accomplished teacher in this instructional context. This is apparent with regard to general organization (the class is structured so that students know at any time what they're supposed to be doing, and what the point of it is) and with regard to the general classroom ambiance (which is supportive of interaction, low-key and informal, but in which the students stay on task). It is also apparent at a more "local" level in specific interactions with students - e.g., his use of reflective tosses is effortless. That he has these skills is important in a number of ways. Obviously, they contribute to his ability to deal easily with routine matters. But, they help with the non-routine as well. Since Minstrell can move things along in the classroom smoothly and easily, he is relatively free to concentrate on difficult issues when they arise.

A third aspect of Minstrell's knowledge is his pedagogical content knowledge, his specific knowledge the ways that students are likely to understand and/or have difficulties with the topic at hand. Minstrell has taught this lesson many times. He knows which issues are likely to be problematic for the students, and he is prepared with examples whose discussion will help clarify the issues. Minstrell knows how to frame his questions so that they lead the class into productive directions. He has a wide range of facilitative moves - classroom routines, easily accessible scripts for making certain points, etc. - that, as in the case with his general pedagogical knowledge, allow him to move comfortably because he is on familiar ground.

It should be stressed that although these categories of knowledge have been named separately, they get accessed and used as part of a seamless whole. To make this point, I return briefly to the notion of interactive elicitation.

A word of preface is necessary before I continue. In what follows, I am about to describe the essence of interactive elicitation as a rather simple decision procedure. It must be understood that no claim is being made that Minstrell follows this decision procedure, either consciously or directly. Nor should it be inferred that a class discussion that has the form of interactive elicitation has any sort of mechanical feel to it (precisely the opposite!). The idea, rather, is that people can act in complete accord with a decision procedure, while not being able to elucidate that procedure - or even being aware that they are acting in accord with it. By analogy, consider the fact that most people are pretty good at turn-taking in conversation. Hardly anyone can describe the rules for turn- taking, which have been elaborated in detail by linguists - and yet people act in accord with those rules nonetheless. That caveat mentioned, a decision procedure that models interactive elicitation is given in Figure 6.

1. Minstrell raises an issue and asks students to say something about it. (For example, "Are there reasons for not considering all the data?")
2. One or more students respond.
3. (Optional) Minstrell may summarize or provide additional detail.
4. If appropriate Minstrell solicits clarifications and elaborations, using questions to shape the direction of the conversation. This continues until the topic has been adequately explored.
5. Minstrell invites a continuation of dialogue: "Anything else we can say? Typically, one of three things occurs:
   a. Another issue is raised, in which case steps 3, 4, and 5 are repeated.
   b. A long silence suggests the end of student contributions, but an important issue has not been discussed. Minstrell raises it for discussion. Steps 3, 4, and 5 are repeated.
   c. A long silence suggests the end of student contributions, and the discussion has covered the expected territory. Minstrell may summarize what the class has discussed (this is optional). He then moves on to the next part of the lesson.

Figure 6.
A simple decision procedure describing interactive elicitation.

It is important to realize that the implementation of interactive elicitation draws upon all aspects of Minstrell's knowledge base, and that in use the distinctions among subject matter knowledge, general pedagogical knowledge, and pedagogical content knowledge become irrelevant. Consider, for example, the sequence that begins with the student's unanticipated suggestion that the class consider her "complicated" formula. Minstrell's immediate recognition that what she proposed can be interpreted as the standard formula for the weighted average is, as discussed above, a matter of subject matter knowledge. His response, which is to have the student clarify her suggestion for the whole class and to get her proposed formula up on the board for discussion, might be considered a matter of general pedagogical knowledge - he is setting the stage for a discussion of her idea. Later on, when the student indicates that she was proposing something different (the denominator in her calculation being the number of coefficients in the numerator, not their sum), Minstrell employs pedagogical content knowledge in moving the discussion to the examination of a

particular example that reveals the properties of this alternative suggestion. In sum: to understand what Minstrell (or any teacher) knows, the different categories of knowledge are important - e.g., having or not having a particular kind of pedagogical content knowledge might be the crucial difference determining the success of an instructional segment. But, when one explores how the knowledge is accessed and used, these categorical differences become irrelevant. The central issue, at a level of mechanism, is the following: given the teacher's beliefs and goals, does the teacher have access to knowledge that can be accessed and used in the service of those beliefs and goals?

With this information as background, I discuss how the lesson actually played out.

6.3. What took place

This section begins with a brief description of the lesson as a whole. It then focuses on the section of the lesson devoted to "best value" (chunk [1.2.2] of Figure 7, below). By way of context, Figure 7 provides the first few levels of the parsing of the whole lesson. What follows is a narrative summary.

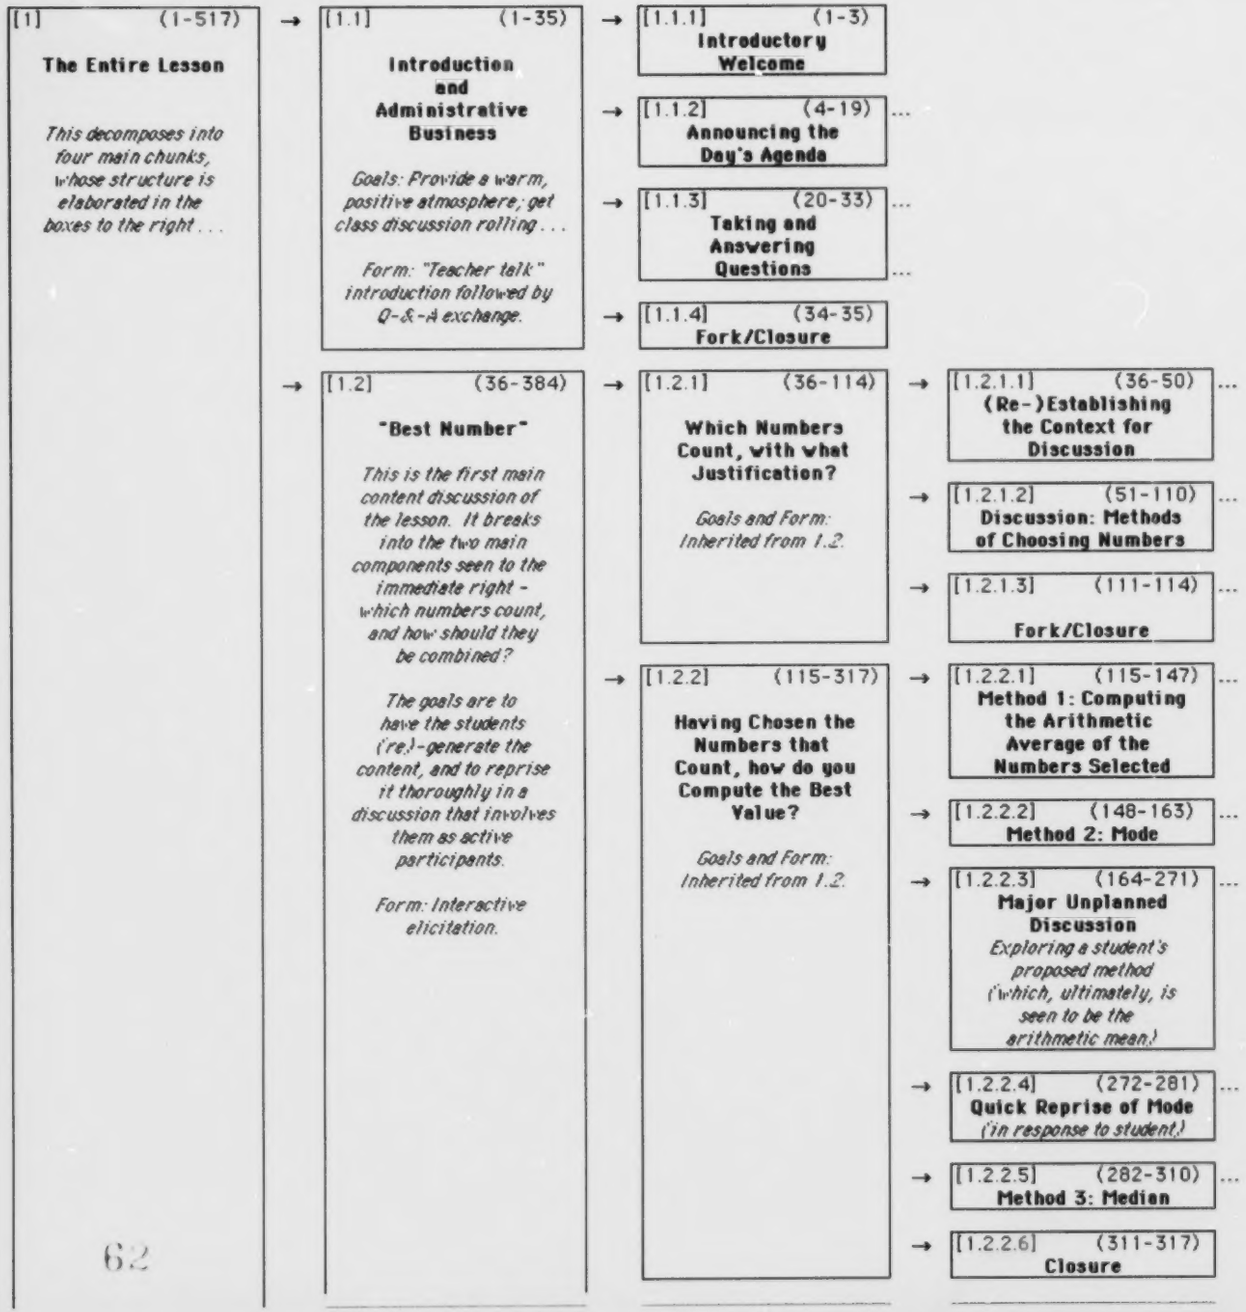
[1] (1-517)
The Entire Lesson
This decomposes into four main chunks, whose structure is elaborated in the boxes to the right . . .
[1.1] (1-35)
Introduction and Administrative Business
Goals: Provide a warm, positive atmosphere; get class discussion rolling . . .
Form: "Teacher talk" introduction followed by Q-&-A exchange.
[1.1.1] (1-3)
Introductory Welcome
[1.1.2] (4-19)
Announcing the Day's Agenda
[1.1.3] (20-33)
Taking and Answering Questions
[1.1.4] (34-35)
Fork/Closure
[1.2] (36-384)
"Best Number"
This is the first main content discussion of the lesson. It breaks into the two main components seen to the immediate right - which numbers count, and how should they be combined?
The goals are to have the students ('re;)-generate the content, and to reprise it thoroughly in a discussion that involves them as active participants.
Form: Interactive elicitation.
[1.2.1] (36-114)
Which Numbers Count, with what Justification?
Goals and Form: Inherited from 1.2.
[1.2.1.1] (36-50)
(Re-)Establishing the Context for Discussion
[1.2.1.2] (51-110)
Discussion: Methods of Choosing Numbers
[1.2.1.3] (111-114)
Fork/Closure
[1.2.2] (115-317)
Having Chosen the Numbers that Count, how do you Compute the Best Value?
Goals and Form: Inherited from 1.2.
[1.2.2.1] (115-147)
Method 1: Computing the Arithmetic Average of the Numbers Selected
[1.2.2.2] (148-163)
Method 2: Mode
[1.2.2.3] (164-271)
Major Unplanned Discussion
Exploring a student's proposed method ('which, ultimately, is seen to be the arithmetic mean.)
[1.2.2.4] (272-281)
Quick Reprise of Mode
('in response to student.)
[1.2.2.5] (282-310)
Method 3: Median
[1.2.2.6] (311-317)
Closure

Goal traces (see below for a legend describing the goals)

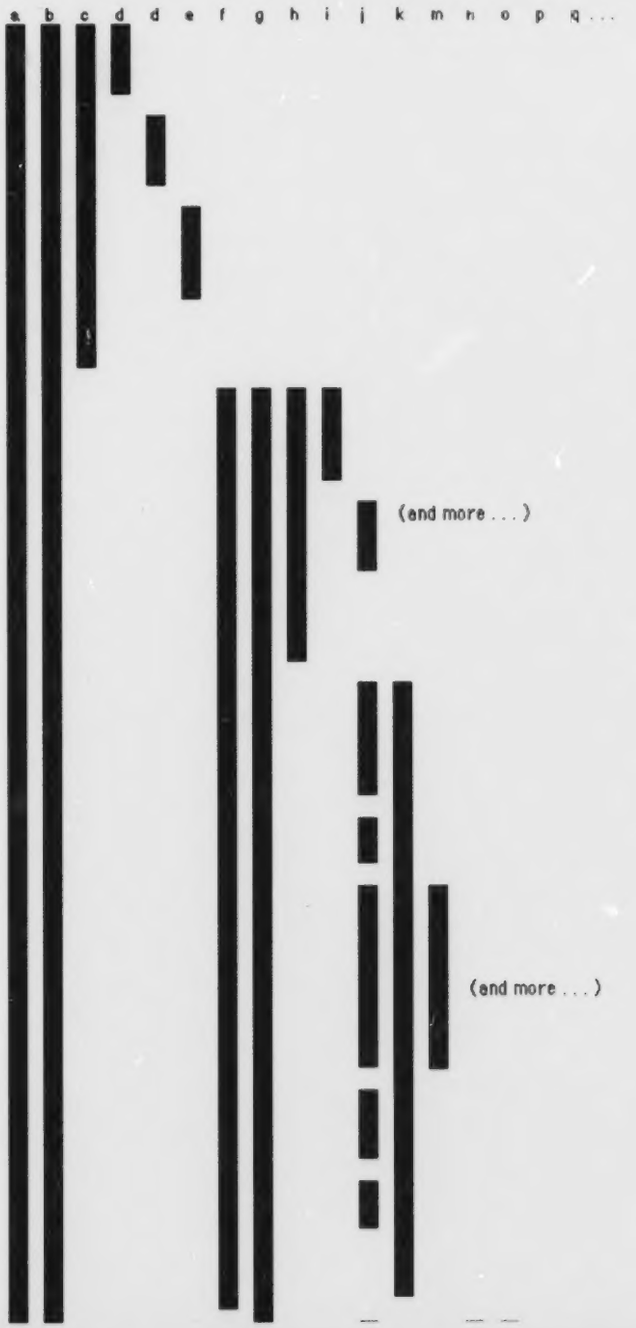

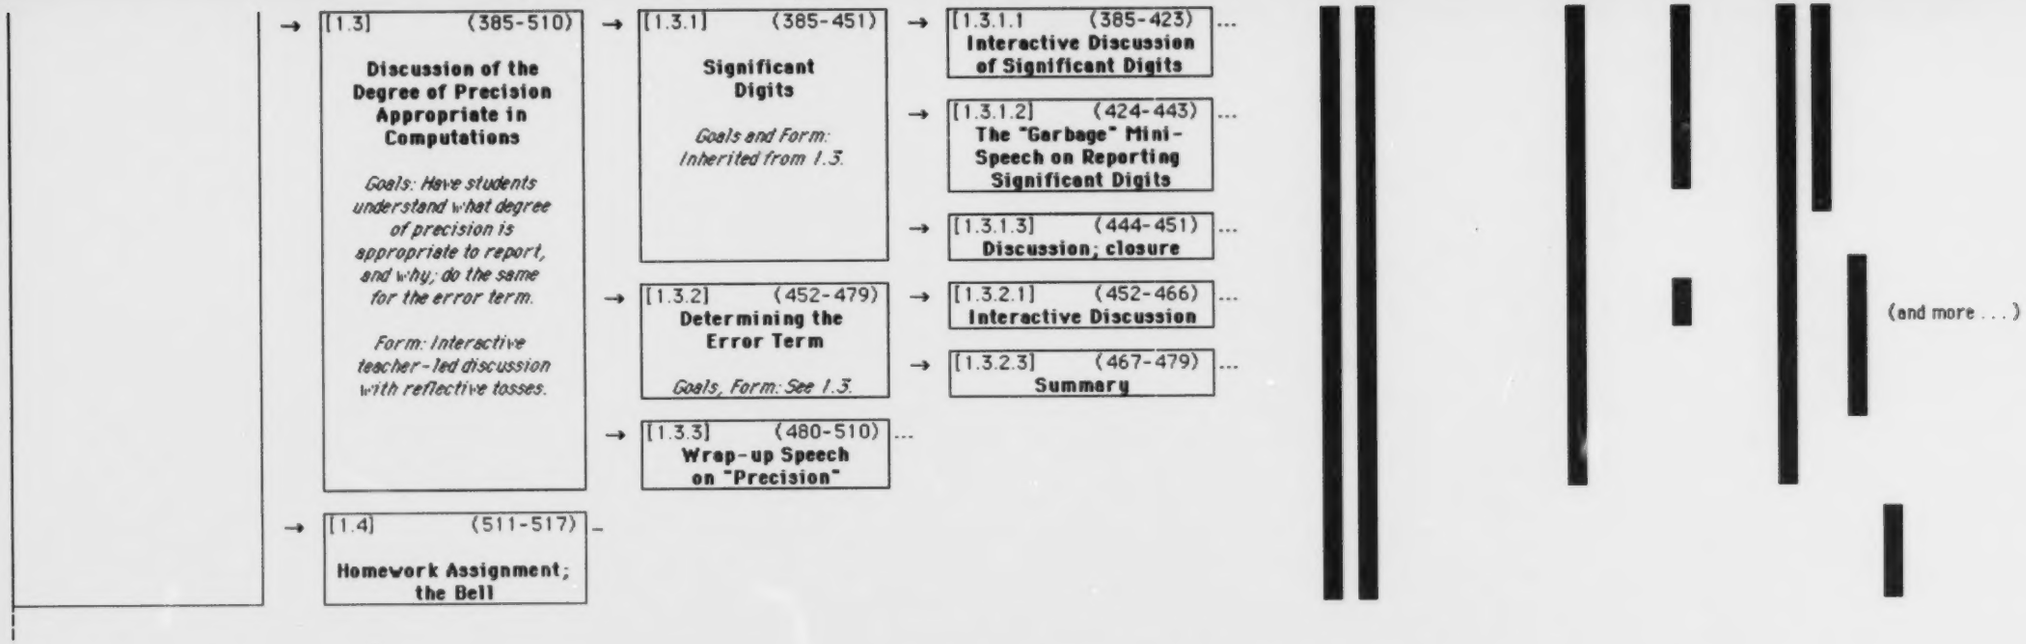
[1.3] (385-510)
Discussion of the Degree of Precision Appropriate in Computations
Goals: Have students understand what degree of precision is appropriate to report, and why; do the same for the error term.
Form: Interactive teacher-led discussion with reflective tosses.
[1.3.1] (385-451)
Significant Digits
Goals and Form: Inherited from 1.3.
[1.3.1.1] (385-423)
Interactive Discussion of Significant Digits
[1.3.1.2] (424-443)
The "Garbage" Mini-Speech on Reporting Significant Digits
[1.3.1.3] (444-451)
Discussion; closure
[1.3.2] (452-479)
Determining the Error Term
Goals, Form: See 1.3.
[1.3.2.1] (452-466)
Interactive Discussion
[1.3.2.3] (467-479)
Summary
[1.3.3] (480-510)
Wrap-up Speech on "Precision"
[1.4] (511-517)
Homework Assignment; the Bell
(and more . . .)

Goal Legend (corresponding to the partial list of goals on the right-hand side above)

a. Have the class interact as a community of inquiry, with freedom to explore, conjecture, reason things through.
b. Have students experience physics as a way of making sense of the world.
c. Provide a warm, positive atmosphere in which students would feel valued, encouraged to speak, etc.
d. Deal with administriva.
e. Provide information about the sequence of the day's activities.
f. Answer any questions the students might have (about anything).
g. Have students conceive of "best number" as a whole - which data, combined in what way?
h. Have students (re)-generate content (via interactive tosses).
i. Flesh out content of "which numbers count."
j. Provide context for the discussion of "best value."
k. Elaborate on specifics of content that arise in dialogue (by interactive elicitation).
l. Have students discuss the three measures of central tendency (mean, median, mode).
m. Pursue student inquiry wherever appropriate.
n. Have students understand "precision": how many digits to report, how to describe the error term.
o. Work through the properties of "significant digits."
p. Work through the details of the error term.
q. Wrap up the class, including assigning homework.

Figure 7. A "shallow parse" (the first few levels of parsing) of the transcript, with a partial goal trace

Overview of the lesson

As indicated by chunk [1.1] and its further decomposition into [1.1.1], [1.1.2], and [1.1.3], the beginning segment of the lesson took place essentially as envisioned. Minstrell took care of administrative business, invited questions and responded to them, and then turned to the main business of the day.

Chunk [1.2] of the lesson played out in the expected spirit but differently in detail. Things began as planned. The discussion of which numbers should be taken into account and why (chunk [1.2.1] and its elaborations) proceeded as expected, as did the transition to the discussion of best value. The discussion of best value [chunk 1.2.2], which will be examined in detail below, contained a major unplanned excursion in response to a student's unexpected suggestion of a "complicated" method for averaging a collection of numbers. By the time that discussion was concluded, however, the intended content had been covered. The next chunk [1.2.3] represents a slight deviation from Minstrell's plan. Minstrell had intended to deal with rounding as part of his "precision" discussion, but a student raised an issue related to rounding during the best value discussion - so Minstrell dealt with it when it occurred. Following that discussion, Minstrell returned to best value and, in chunk [1.2.4], brought that part of the lesson to a close.

Chunk [1.3], devoted to significant digits and measurement error, took place as planned - although it was somewhat condensed because of the time that had been spent on the unplanned excursion in [1.2.2]. Minstrell finished this discussion as the bell rang, and quickly (chunk [1.4]) assigned homework for the next day.

Detail: Best value

Here I focus on chunk [1.2.2], with an emphasis on [1.2.2.3], in which Minstrell made an impromptu departure from his lesson plan to pursue an issue raised by a student. Appendix B provides the transcript of [1.2.2], which comprises lines 115-317 of the transcript of the whole lesson. Figure 8, which continues the parsing given in Figure 7, provides a more detailed analysis of [1.2.2.3].

Here is a narrative version of what took place in [1.2.2]. Things began as planned, with chunks [1.2.2.1] (lines 115-147) and [1.2.2.2] (lines 148-163) playing out in precise accord with Minstrell's lesson image. Indeed, these lesson segments provide textbook examples of Minstrell's questioning methods. Minstrell opens the discussion in lines 115-120 with a general question, "What's one thing that we might do with the numbers?" When a student replies "average," he asks for a clarification (lines 122-123). The student provides the desired information (lines 124-125), at which point Minstrell provides some elaboration (lines 126-139). This brings the discussion of average, and [1.2.2.1], to a close. Minstrell then asks the generic transition question, "Any other suggestions . . . ?" (lines 140-142). To give the students time to generate an answer and to provide more grist for the conversation, he puts on the board the data the students had obtained for the table's length and width. The numbers that become the focus of the conversation are the values obtained by the students for the table's width in centimeters,

106.8; 107.0; 107.0; 107.5; 107.0; 107.0; 106.5; 106.0.

A student observes (line 148) that one particular number, 107.0, appears repeatedly. Minstrell asks for a clarification, and then in response to the student's comment, points toward the definition with another question: "Is there any other number in the width column that shows up as much as 107 point zero?" When the answer is "No," he reminds them (lines 155-159) that the number that appears most frequently is called the mode, and that mode is "another way of getting a best value out of a collection of numbers that you're willing to keep." This brings [1.2.2.2] to the point of closure. Minstrell does a quick status check (lines 160-162) and then returns (line 163) to his generic prompt: "Anybody think of another way of giving a best value?"

At this point a student makes the following suggestion (lines 165-167), which initiates [1.2.2.3]:

> "This is a little complicated but I mean it might work. If you see that 107 shows up four times, you give it a coefficient of 4, and then 107.5 only shows up one time, you give it a coefficient of 1, you add all those up and then you divide by the number of coefficients you have."

In lines 168-175, Minstrell's requests for clarifications result in the public elaboration of one possible meaning of the student's suggestion. In interaction with the student he produces this version of the formula the student had suggested:

$$\frac{4(107.0) + 1(106.8) + 1(107.5) + 1(106.5) + 1(106.0)}{8}.$$

He then asks the students what they think. Among the clearly audible comments is "It's the same," which Minstrell once again turns back to the students as a question (lines 204-5): "All right. So actually it ends up being the same as the arithmetic average?"

Interestingly, the student who first proposed this formula, S8, replies in the negative (line 206): "No. Because 107 gets four times the value, so the 107 counts more." In a subtle and powerful move, Minstrell does not respond directly but turns to the student who had initially (in line 121) defined "average," and, in lines 207-210, asks that student for an elaboration of her original definition. They work out the implications of the definition together in lines 211-225, by which time the whole class agrees that the two methods (arithmetic average and the formula as interpreted above) are the same.

Agreeing with this result, S8 notes that she didn't realize it when she proposed the example. (Recall her comment about 107 "counting more.") This raises (to Minstrell) the question of comparing the weighted and the unweighted average, so he poses it to the class (lines 227-235). When the class asserts that the two formulas would produce different results, he pushes further and asks if they can provide a "real clear example" of why the two would differ. He then pursues and (through line 271) elaborates an example proposed by a student. This concludes the unplanned excursion.

The rest of chunk [1.2.2] moves briskly. After another status check in line 272 ("Anybody confused yet?") Minstrell returns to his generic question: "All right. Anybody else see a different way of approaching?" In the context of the example they had just worked through a student suggests that the best value might be the number that appears most frequently. Minstrell reminds the class that this is the mode, bringing this short chunk ([1.2.2.4], lines 271-281) to a close. Then since the class has failed to mention the median, he brings it up for discussion in lines 282-285. Again, he does so with a question: "Anybody know what the median is?" In lines 288-309 he works through the definition of median, and in lines 311- 317 he brings this part of the lesson to a close with a summary of the three main ways one might combine the data (arithmetic average, mode, and median).

6.4. Analysis

In parallel with the analysis given in Section 5.4, I examine the ways in which Minstrell's beliefs, goals, lesson image, and knowledge base shaped what took place in the classroom. To begin, I note that Minstrell has taught, reflected on, and revised this instruction for some years. He is articulate about his beliefs and goals, which are quite consistent. This lesson (and thus the lesson image corresponding to it) was constructed explicitly with specific high priority goals in mind. Moreover, Minstrell has a knowledge base that enables him to implement

the lesson in a way consistent with his lesson image, but with flexible adaptations when contingencies arise.

As in the preceding section, my comments proceed at two levels: first with regard to the lesson as a whole, second with regard to the unplanned excursion in chunk [1.2.2].

The lesson as a whole

It is worth noting that, even with the substantial alteration made in chunk [1.2.2], the overall structure and character of the lesson conform in large measure to Minstrell's lesson image.

First, structure. The introduction, [1.1], goes as planned. At the macro level, the same can almost be said of the discussion of best value in [1.2]. That discussion has two main parts, [1.2.1] and [1.2.2], which deal respectively with issues of data selection and how to combine the selected data; there is also a summary wrap-up in [1.2.4]. The only diversion from Minstrell's expected content coverage is the discussion of rounding in [1.2.3], which takes place earlier than planned because a relevant issue was raised by a student (its expected place was in [1.3]). The balance of the lesson, in chunks [1.3] and [1.4], deals with the expected material, though at a faster clip than Minstrell had originally envisioned.

Second, character. In a global sense, the lesson is entirely consistent with Minstrell's high priority goal of fostering a community of sense-making in which exploring ideas is highly valued. (Indeed, the excursion in [1.2.2.3] underscores these commitments.) Moreover, the general instructional character of each segment conforms to Minstrell's questioning style (interactive elicitation) as described in Section 6.2. For example, the discussion of which numbers one might select, in [1.2.1], takes place exactly as envisioned: Minstrell asks for suggestions from students, clarifies the suggestions by use of reflective tosses, asks for more suggestions, clarifies those, and continues in that fashion until the well runs dry. Likewise, even though there is an unplanned excursion in [1.2.2.3], chunk [1.2.2] proceeds in precisely the same manner. Minstrell opens the floor for comments, and when a student suggests "average them" he uses reflective tosses to clarify the definition of arithmetic average (chunk [1.2.2.1]). He asks if there are other methods, and in response to a student comment that one number appears more frequently than others, clarifies the definition of mode (chunk [1.2.2.2]). He asks for more comments, and a student suggests the "complicated" formula whose properties he winds up exploring carefully (chunk [1.2.2.3], further analyzed below). He asks again, gets a comment about mode, and then, since the well has now run dry, introduces median and clarifies its meaning by way of reflective tosses. This conforms precisely to his style and lesson image. Finally, chunk [1.3], though done at a pretty fast clip, is also consistent with Minstrell's intentions.

These observations raise some interesting questions. As noted, the lesson took an extended detour in response to a comment made by a student. Even with this detour, the lesson as a whole appears natural and unforced. Yet, Minstrell managed to do many of the things he wanted to do. How was he able to do so?

A close analysis suggests that the answer resides in the collection of resources Minstrell brings to the lesson (in the context of his beliefs and goals, of course). As suggested in the discussion of Minstrell's knowledge base in Section 6.2, Minstrell's interactions with the students shape and contribute to the ongoing dialogue in subtle ways, serving to move it in very particular directions. Some of his questions simply ask for clarification, and if the clarification is forthcoming, no more is needed - see, e.g., the exchange in lines 122-125 of Appendix B. But in other cases, he takes the lead. For example, when a student notes that "You've got a bunch of numbers that are the same number" (line 148), Minstrell's first response is generic: "OK. Like what are you talking about there?" This question leaves lots of room for response. Given where Minstrell wants to go, the student's response, "107," is correct but not especially helpful. Minstrell's response is interesting. He validates the studen*'s response ("All right. 107 point zero, 107 point zero, 107 point zero, 107 point zero"), but then goes on to ask a question that points to the critical part of the definition of mode: "Is there any other number in the width column that shows up as much as 107 point zero?" In this way he provides positive feedback and moves the agenda forward, while still asking questions rather than simply "telling."

For the most dramatic example of how Minstrell uses what he knows to foster student contributions to the discussion, see the discussion of [1.2.2.3] that follows. I shall return to the issue of the knowledge base in the concluding discussion.

Detail: Best value

Issues of beliefs, goals, and resources are central to understanding what takes place in [1.2.2.3]. I begin with beliefs. Central to Minstrell's conception of teaching is the belief that student contributions should be encouraged and pursued, where he sees them as being relevant and constraints make it possible to pursue them. In addition, Minstrell believes that doing physics is a sense- making activity, and he wants his students to experience it as such. Hence, when questions about the properties of objects they are studying come up in the class, Minstrell will want to have the class sort things out. Given these beliefs, Minstrell's decisions in [1.2.2.3] are exactly what you would expect. When S8 makes a suggestion regarding a possible formula for computing the average, pursuing it becomes a high priority goal for Minstrell. Likewise, when S8 comments that she was unaware that her suggested formula was the same as the average, clarifying the relationship between the two different forms of average is assigned very high priority and moves to the top of Minstrell's goal stack. In

focusing on these issues, Minstrell honors student inquiry and fosters sense- making at the same time.[21]

How Minstrell chooses to pursue these issues is worthy of note, for it points to the resources at his disposal for the implementation of his pedagogical vision. As noted above, Minstrell's preferred form of interaction is the reflective toss. In response to S8's suggestion, he first asks for clarification, and then, in interaction with her, writes her proposed formula on the board. This sequence of actions makes her suggestion public, and the property of the whole class. He then asks the class what it thinks of the proposed method. The primary responses are "too hard," which he echoes with a question ("Too hard?"); the responses to this include at least one comment that, "It's the same" and a comment from S8, "No. Because 107 gets four times the value, so the 107 counts more." Now, how does Minstrell pursue the issue that has suddenly emerged with S8's comment? He calls on S5, who had first suggested arithmetic average, and the following dialogue ensues:

S5: You'd add all the numbers together and then divide it by 8.

T: Now what do you mean by "adding all the numbers"?

S5: You would add each separate number that everybody got; you wouldn't just add one 107, you'd add all the 107s.

By virtue of this choice, he insures that the main content contribution comes from a student.

As we leave this example, I want to stress once again the role that Minstrell's knowledge base plays in what he is able to do. As discussed in Section 6.2, his almost effortless-seeming response to this unexpected student contribution rests heavily on Minstrell's subject matter knowledge, general pedagogical knowledge, and pedagogical content knowledge. The fact that he immediately recognizes one possible interpretation of the student's suggestion to represent the weighted average means that he is free to think on his feet - he does not have spend time sorting out for himself what the suggestion might mean. That is, knowing the territory means that he can be a guide to it. Moreover, his techniques for guiding students through the territory are well honed. The subject matter content and the questioning style he uses are familiar. This puts him in a position where he is able to frame questions that build on and clarify student contributions, shape the direction of inquiry, and make substantive contributions to the conversation itself.

---

[21] A brief representational note. One thing the parsing scheme does not do especially well is highlight the relationships between constellations of beliefs and the goals that correspond to them. For example, the story I just told regarding Minstrell's unplanned excursion in Chunk [1.2.2.3] can be found in the first box of Figure 7 - but it hardly jumps out at the reader.

Beliefs, goals, and knowledge combined - the issue of mechanism

In the preceding discussions I have focused on the individual components of the theory and model - on Minstrell's beliefs, his goals, his lesson image, and the mental resources (the knowledge base) that he has available. Minstrell has a strong and well articulated belief system regarding his sense of the discipline (physics as a sense-making activity) and how students should learn it (as members of an intellectual community in which thoughtful exploration and justification are highly valued). He has equally strong and closely linked goals for instruction (see, e.g., van Zee & Minstrell, 1997b). And, as summarized in the previous paragraph, his knowledge base provides substantial resources to be used in the pursuit of those goals. Now, having described the components, I turn to issues of mechanism: how do the pieces fit together?

The beginning of chunk [1.2.2.3] provides the perfect case example for a discussion of the interactions of beliefs, goals, and knowledge at a level of mechanism. Let us begin at line 163 of Appendix B. To this point chunk [1.2.2] has gone exactly as envisioned, with discussions of average (in [1.2.2.1]) and mode (in [1.2.2.2]) conducted via interactive elicitation. High priority goals concern the continuation of this discussion, with the expectation that median will be the next topic covered. The completion of [1.2.2.2] corresponds to step 4 in the model of interactive elicitation described in Figure 6. This serves as a trigger for step 5, in which (according to the model) Minstrell will invite a continuation of dialogue. Indeed, he does, in line 163: "Anybody think of another way of giving best value?"

Now, S8's response, "This is a little complicated . . . " is not what Minstrell expects. How would we expect him to react, and why? Consider the following. First, Minstrell's epistemological stance includes the belief that student inquiry should be honored and encouraged, if possible and appropriate. Second, Minstrell immediately recognizes the suggestion as providing a potential lead into important (though unplanned) content - the formula for weighted average, which is indeed relevant and appropriate. It is still early in the lesson, so there is time to pursue this opportunity if he wishes to. Will he do so? Given this characterization of his beliefs and knowledge, the answer has to be yes. In Minstrell's view the student's suggestion is relevant and appropriate; there is time to pursue it; doing so will lead to an enriched discussion of content; doing so will clearly reinforce Minstrell's message that student suggestions will be taken seriously and may well result in contributions to everybody's knowledge. So, Minstrell will establish the goal of following up on the student's comment. Now, how will he do so? By reflective tosses, the method of choice in his pedagogical tool kit. (In terms of the model, the use of reflective tosses/interactive elicitation is the action plan that Minstrell selects to pursue the new high priority goal of working through the student's suggestion.)

From this point on the implementation is straightforward. Since the student's suggestion went by rather fast, Minstrell's first move is to make sure that the class understands what is to be considered. His move "You lost me" in line 168 is an invitation for clarification, which takes place in lines 170 through line 199. Then what? The fact that the problem is now well framed serves as a trigger for the next step: a reflective toss asking the students what they think.

If space permitted I could continue at this level of description. Things proceed as planned through line 225, where Minstrell does a status check (part of his routine). The comment from S8 in line 226 that she hadn't realized that her formula was the same as average raised further content-related issues for Minstrell, for lurking behind the student's comment was the fact that there are ambiguities regarding the ways numbers "count" in formulas with coefficients. Clarifying these potential ambiguities becomes a high priority goal. Minstrell does so via his questioning strategy, in lines 227-271. At that point, this episode is effectively completed. This corresponds to the completion step 4 in the description of interactive elicitation (recall Figure 6). He then moves to step 5 in line 274, "all right. Anybody else see a different way of approaching?"

Let me now characterize the previous description in abstract terms. The student's suggestion in line 164 served as a trigger for the very high activation of particular beliefs (specifically, related to honoring and encouraging student inquiry). An exploration of the teacher's knowledge base, and a constraint check, suggested that the issues raised by the student could be pursued profitably. Hence a new goal - to work through the student's suggestion - was given highest priority. The teacher then selected and implemented an action plan (to use reflective tosses to clarify the student's suggestion and pursue its meaning) that was consistent with this new constellation of high activation beliefs and goals. This action plan was carried out in lines 168-225. Lines 221-225 were a "status check," the last part of the action plan. Had no new issues arisen, the deferred goal of continuing "best value" via interactive elicitation would have returned to high priority, and the discussion of best value would have continued at that point. However, a new issue did arise in line 226. This triggered the activation of a new high priority goal, to clarify the difference between weighted and unweighted averages. Once again, an action plan consonant with the current configuration of high activation goals and beliefs was selected and implemented. A status check at the conclusion of that action plan (lines 272-3) indicated that the relevant goals had been attained. That served as a trigger for the goal that had been put temporarily on hold - continuing the "best value" discussion via interactive elicitation - to return to highest priority. The lesson then continued as (originally) planned.

6.5. Could we model this lesson in detail?

In a word, yes - and the discussion in the previous paragraphs indicates precisely how and why. Though decidedly more complex than Nelson's lesson in many ways, Minstrell's lesson is in fact easy to model. Minstrell's lesson image, which provides the base for modeling the lesson, is rich in detail. It provides the substance for navigating familiar terrain smoothly. (And, I note, most of the relevant terrain *is* familiar.) Minstrell's beliefs are clear and well articulated, so predicting the decisions he might make in somewhat unusual circumstances (e.g., whether or not he will choose to pursue the issue raised by S8 in line 164) is straightforward. Moreover, when Minstrell does decide to follow a particular path, the means by which he will pursue it - his questioning strategies - are also fairly easy to model. As the discussion in Section 6.4 indicates, it is a relatively straightforward matter to model and even to predict, at a level of mechanism, what Minstrell does and why.

There are, of course, limitations to what can be predicted. Some of what Minstrell says or does is ad hoc (see, e.g., the discussion of "operational definition" in lines 126-129 of the transcript), and there is no way of knowing whether he would say or do similar things in similar circumstances. Elsewhere, some of the things Minstrell does fall within a constrained range of possibilities, and there is no way of knowing which of them he might choose. For example, the model predicts that he will respond to S7's comment in line 148, "You've got a bunch of numbers that are the same number," with a reflective toss that leads to the definition of mode. The precise way he will do so - in this case by asking whether any numbers show up as frequently as 107.0 - is beyond the realm of prediction. But these are low-level details. For the most part, Minstrell's classroom decisions and actions - which are anything but simple or rule-bound - can be modeled with great fidelity and without much difficulty.[22]

[22]The discussion in Section 6.4 suggests how much of the modeling is done, at a level of mechanism. For more detail see Schoenfeld, Minstrell, & van Zee, 1996.

7. CASE STUDIES 3 AND 4 IN BRIEF: ALAN SCHOENFELD AND DEBORAH BALL

This section of the paper describes two additional cases. The first provides additional evidence of the robustness of the model, while the second points to some possible limitations in terms of the model's scope. I provide very brief descriptions of the cases here. Their implications are pursued in Section 8.

7.1. Case 3: The opening days of my problem solving course

The third case is the study of my own teaching, specifically the first few days of the course in mathematical problem solving I offer for undergraduates at Berkeley.[23] Despite some non-trivial differences in level and content (university- level mathematics versus high school physics), there are very strong parallels between these lessons and the lesson taught by Minstrell that was discussed in Section 6. My course, like Minstrell's, was designed to introduce students to thinking within a particular scientific discipline as a form of sense-making. The course has been the object of reflection and refinement for many years, so most of what takes place in it, especially during the first few days, is well practiced. I have clearly elaborated goals for the course as a whole and for the beginning lessons in particular. I also have, as delineated in Section 3.2 of this paper, a well fleshed out lesson image (which includes what I plan on doing, how I expect the students to respond, and how I expect to respond to their responses) and well- rehearsed action plans to be employed in the service of those goals.

My first main purpose in discussing the opening lessons of the problem solving course is to bring some closure to the discussion of lesson image and action plans given in Section 3.2. The second is to indicate that, just as in the two previous case studies, constructing a detailed model of this chunk of instruction would be a rather straightforward matter. The evidence offered here is suggestive only. For extensive detail, see Arcavi, Kessel, Meira and Smith (1998) and Schoenfeld (1998).

The lesson as a whole

Figure 9 provides the first few levels of detail of the parsing of the first lesson of the 1990 version of the problem solving course, which is the version analyzed by Arcavi et al. The main thing to observe - no great surprise - is the correspondence, at a global level, between the structure of the lesson as it unfolded and the lesson image for that instruction described in Section 3.2.

---

[23] I should note that I originally had no intention of modeling my own instruction. In a separate study, however, Arcavi, Kessel, Meira, and Smith (1998) did an extensive analysis of the first few days of the 1990 version of the course. With that much analysis having been done independently, it seemed appropriate to use their data to build a model of the instruction, and to see what "value added" there might be to the model. See Schoenfeld (1998) for more detail.

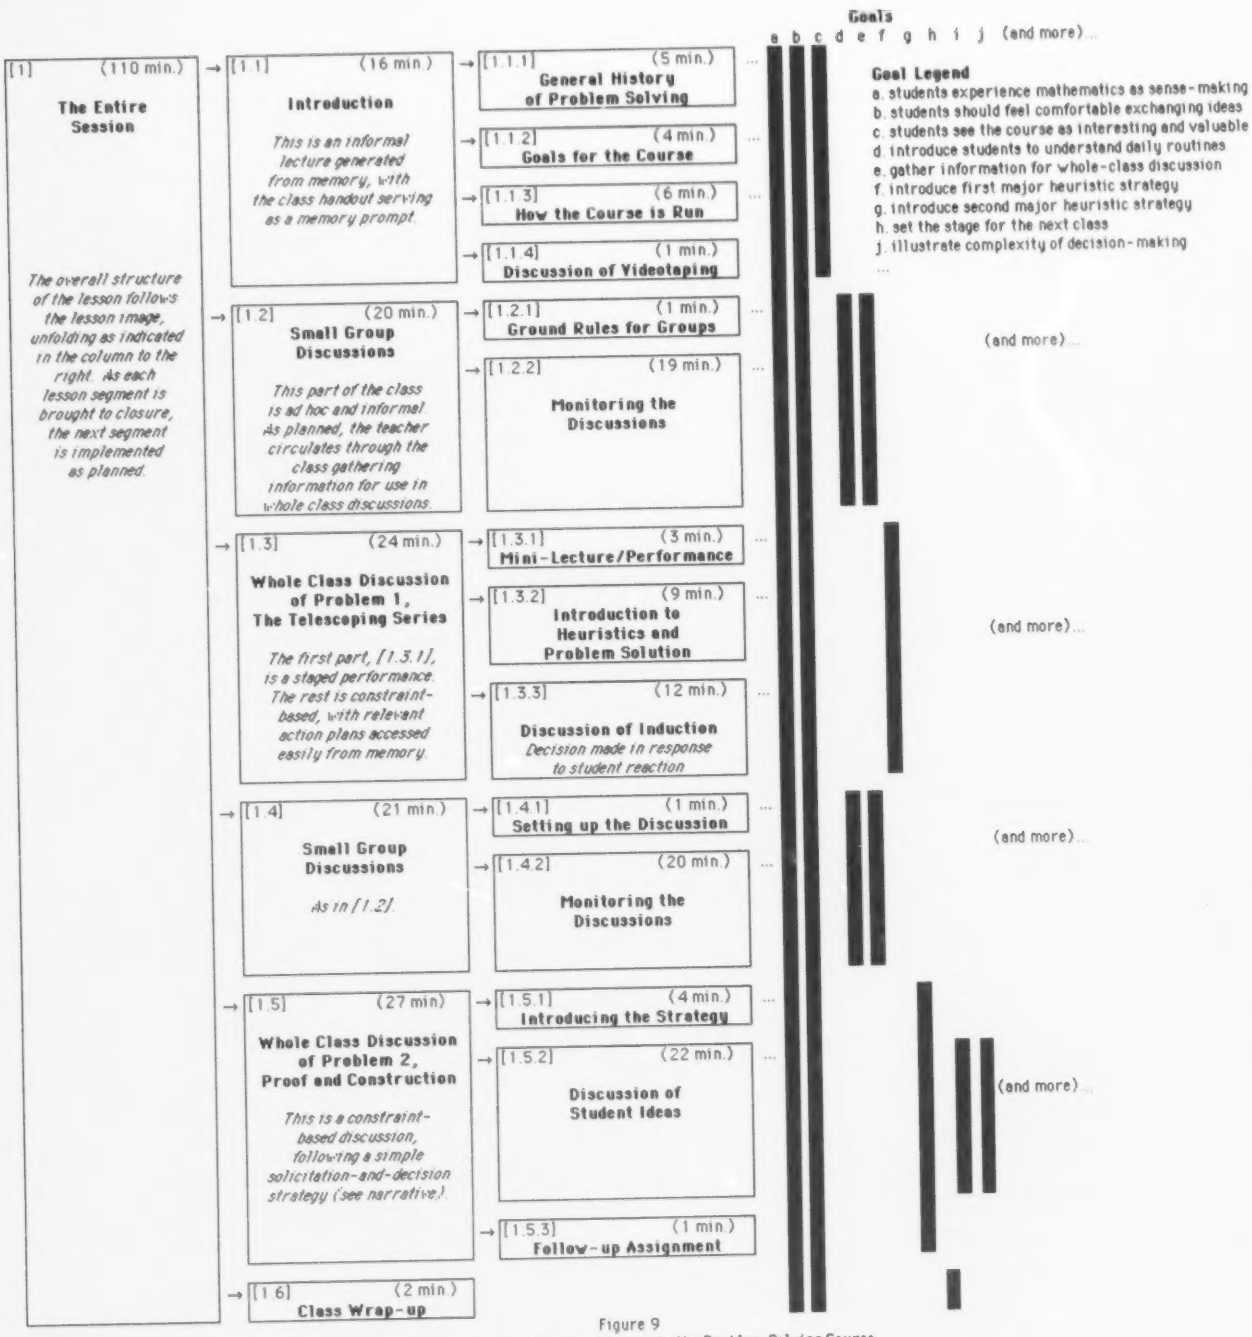
Goals
a b c d e f g h i j (and more)...
[1] (110 min.)
The Entire Session
The overall structure of the lesson follows the lesson image, unfolding as indicated in the column to the right. As each lesson segment is brought to closure, the next segment is implemented as planned.
[1.1] (16 min.)
Introduction
This is an informal lecture generated from memory, with the class handout serving as a memory prompt.
[1.1.1] (5 min.)
General History of Problem Solving
[1.1.2] (4 min.)
Goals for the Course
[1.1.3] (6 min.)
How the Course is Run
[1.1.4] (1 min.)
Discussion of Videotaping
[1.2] (20 min.)
Small Group Discussions
This part of the class is ad hoc and informal. As planned, the teacher circulates through the class gathering information for use in whole class discussions.
[1.2.1] (1 min.)
Ground Rules for Groups
[1.2.2] (19 min.)
Monitoring the Discussions
[1.3] (24 min.)
Whole Class Discussion of Problem 1, The Telescoping Series
The first part, [1.3.1], is a staged performance. The rest is constraint-based, with relevant action plans accessed easily from memory.
[1.3.1] (3 min.)
Mini-Lecture/Performance
[1.3.2] (9 min.)
Introduction to Heuristics and Problem Solution
[1.3.3] (12 min.)
Discussion of Induction
Decision made in response to student reaction
[1.4] (21 min.)
Small Group Discussions
As in [1.2].
[1.4.1] (1 min.)
Setting up the Discussion
[1.4.2] (20 min.)
Monitoring the Discussions
[1.5] (27 min)
Whole Class Discussion of Problem 2, Proof and Construction
This is a constraint-based discussion, following a simple solicitation-and-decision strategy (see narrative).
[1.5.1] (4 min.)
Introducing the Strategy
[1.5.2] (22 min.)
Discussion of Student Ideas
[1.5.3] (1 min.)
Follow-up Assignment
[1.6] (2 min.)
Class Wrap-up
(and more)...
(and more)...
(and more)...
(and more)...
Goal Legend
a. students experience mathematics as sense-making
b. students should feel comfortable exchanging ideas
c. students see the course as interesting and valuable
d. introduce students to understand daily routines
e. gather information for whole-class discussion
f. introduce first major heuristic strategy
g. introduce second major heuristic strategy
h. set the stage for the next class
j. illustrate complexity of decision-making
...

Figure 9
A Top-Level Parse of the First Lesson in the Problem Solving Course

BEST COPY AVAILABLE

The introduction took place more or less as envisioned, as did the transition to small-group work. The work in small groups proceeded as planned, modulo one judgment call. In the past the course had met for three hours per week in two 90-minute periods. A departmental reorganization had resulted in the course being offered for four units, in two two-hour meetings. During the 90-minute periods, I had typically had the students engage in one session of group work followed by whole-class discussions. This time, however, I decided (on the spot) that the two-hour class period would require too long an uninterrupted period of small- group work - especially on the first day, when they were unused to such activity. After about twenty minutes I convened the class for a whole-class discussion of the first problem.

The ensuing discussion, which took about 24 minutes, served multiple goals. As planned, it provided the opportunity for my mock lecture caricaturing the "typical calculus professor" and showing the students that I knew the mathematics and had a good sense of their experience in typical mathematics courses. (This, along with parts of the introductory lecture, plays an important role in helping to convince the students that this decidedly non-standard course may be worth staying in.) It also served as the substantive introduction to our first problem solving strategy (which enabled the students to solve in a minute or two a problem that they had just been unable to solve) and for a review of induction. But, it is worth noting that the timing of the intervention was also important. Originally, my plan had been to extend the small-group part of the class to a length that would accommodate the new 120-minute class length. However, my sense of the class after about 20 minutes of small-group work was that they were beginning to feel slightly uneasy, even though I had been making various suggestions to the groups as they worked on the problems. Their model of instruction was that the teacher's job is to tell them things. Had I abdicated that responsibility? Were they going to be on their own, solving problems without guidance or discussion? The emergent goals of providing reassurance and showing that even though this was a non-standard class, they would be getting some of what they expected, resulted in my bringing the small group work to a close and moving to whole-class discussion. Once the extended discussion of the first problem was completed, I felt comfortable moving back to small groups. At the macro level, the rest of the lesson evolved in concert with my lesson image.

Similar comments could be made at the next finer levels of grain size, which corresponds to the third column of Figure 9 (and beyond). Generally speaking, my plan is constraint-based: I will ask students for suggestions about how to approach the problems, and then work with what they provide.

Based on my experience with these particular problems, I can anticipate most of the suggestions students will make (although not in what order, of course). There are, as was the case with Minstrell, fail-safe strategies in case the students

do not suggest what I expect them to. More fine-grained decisions are often made on-line, in an ad hoc manner.

An example in some detail: The discussion of one problem solution

One of the problems we discuss the first day of the course is a geometric construction problem modified from Pólya (1945) as follows:

> You are given the triangle on the left in the figure below. A friend of mine claims that she can inscribe a square in the triangle - that is, that she can find a construction, using straightedge and compass, that results in a square, all four of whose corners lie on the sides of the triangle. Is there such a construction - or might it be impossible? Do you know for certain there's an inscribed square? Do you know for certain there's a construction that will produce it?

The given triangle

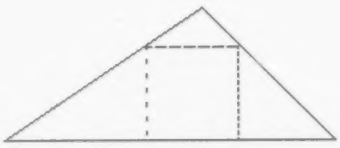

What you'd like to get

The problem serves as an introduction to the following general heuristic strategy: *If you can't solve the given problem, try to solve an easier related problem.* The statement of this strategy is deceptively simple, and its implementation in this problem is actually rather complex. One reason for the complexity is that it is by no means easy for a student to pick the "right" easier related problem. There are many related problems one might consider, for example:

* inscribing the square in a "special" triangle, such as an isosceles or equilateral;
  inscribing a circle in the given triangle;
* constructing a square that has three of its four vertices on the sides of the
* triangle;
  inscribing a rectangle inside the given triangle;
* constructing a triangle similar to the given triangle around a square.
*

As it happens, some of these approaches lead nowhere; some yield existence proofs that there is a square whose four corners lie on the triangle but fail to lead to a construction that can be done with straightedge and compass; and some ultimately lead to a solution of the original problem. An important part of my intended classroom discussion is to wade through this complexity - to have the

students understand that a good deal of discretion is involved in employing these ostensibly simple strategies. They need to learn that weighing various options, selecting among them carefully, and monitoring and assessing one's progress toward a solution are all important parts of problem solving.

In the lesson analyzed by Arcavi, Kessel, Meira and Smith (1998), I posed the "inscribed square" question and asked students to identify the easier related problems that they had considered when they had worked in small groups. As the students generated the approaches, I wrote them on the board. One of the suggestions made by a student was non-standard, and discussing it at that point would have led us temporarily astray. I wrote that suggestion on the board, but apart from the others. In addition, an approach that is often mentioned by students and that is useful to discuss was not mentioned by this group. I brought up the strategy, noting that students in previous versions of the course had raised it and that it was worth considering alongside the ones they had mentioned. When we worked through the various possibilities, I did not take them simply in the order that they had been generated, but in an order that allowed various mathematical "lessons" to emerge more naturally from the discussions. For example, the non-standard suggestion was dealt with last. By the time we got to it, our discussions had provided enough information to render that approach untenable. In consequence, the discussion did not disrupt the flow of the lesson.

This is not much detail, but there is enough here to establish some parallels between the flow of this lesson and that of Minstrell's, discussed above. On the one hand, the interactions were casual and somewhat free-form; I certainly was not following any set of rules or algorithms in deciding what to do as I taught. On the other hand, my on-line decision-making can be modeled with the rather straightforward decision procedure given in Figure 10.

1. Ask for a suggested approach.
2. When a student makes a suggestion . . .
   a. and a discussion of it would be unproblematic, add the suggested approach to the list. Ask if there are more suggestions.
   b. and a discussion of it would be problematic (e.g., an elegant solution resulting from the suggestion might make the discussion of other more cumbersome approaches seem anticlimactic) put that suggested approach on hold, to be discussed after others have been dealt with. Ask if there are more suggestions.
3. If the students fail to make a suggestion . . .
   a. but the list of ideas is relatively complete, then begin a discussion of the ideas in order.
   b. but an important suggestion is not on the list, mention it as one that students have generated in the past, add it to the list, and begin a discussion of the ideas in order.

Figure 10.
A simple decision procedure that models my on-line decision-making in the discussion of the "inscribed square" problem

I will pursue the parallels with Figure 6 in the concluding discussion.

7.2. Case 4: Deborah Ball's "Shea number" class

This discussion of Deborah Ball's "Shea number" lesson,[24] unlike the other discussions in this paper, is unaccompanied by a parsing diagram. The reason, simply put, is that the Teacher Model Group has not yet produced an analysis of that lesson that we find sufficiently compelling to warrant summary in that form.[25] What follows is a brief description of aspects of the lesson and of the issues the lesson raises for us.

---

[24]The lesson, taught January 19, 1990, is part of the data collected by the "Mathematics and Teaching through Hypermedia" project, for which Deborah Ball and Magdalene Lampert were principal investigators. A main purpose of that project is to create a set of common artifacts for discussion in the field - videotapes of teaching and collections of supporting documentation (lesson logs, lesson plans, etc.) that can be examined by teachers, practitioners, and researchers from different points of view. See, e.g., Ball & Lampert, 1996.

[25]Recall that a parsing diagram represents the mere tip of the analytic iceberg. A full analysis, of which the parsing diagram is an iconic summary, includes the attribution of goals and knowledge states that correspond to each chunk of the lesson represented in the parsing.

A number of features distinguish this lesson from the ones discussed earlier in this paper. The students in Ball's class were in the third grade, while those in Nelson's, Minstrell's, and mine were in high school and college. Nelson's lesson was essentially self-contained and was very strongly focused. Minstrell's and mine took place at the beginning of the school year. Both were also strongly focused: Minstrell's was a review-and-elaboration lesson, and mine focused on the introduction of specific new material. In both cases, a major pedagogical goal was the establishment of classroom norms, and a major content-related goal was providing students with an introduction to the discipline as a form of sense- making. In neither case was there much "history" between teacher and the class.

In contrast, the "Shea number" class took place in the middle of the school year. The teacher and students knew each other well, and the social and intellectual norms of the community were well established. The class had a shared language that included *conjecturing* , *disagreeing* with a conjecture, and *revising* one's previous opinion; they understood that when you disagree with something someone has said, you will be expected to provide a reason or evidence for your contrary opinion. This class was one of an extended series of lessons, so it is essential to understand the instructional history (what issues were hanging in the air unresolved, and what the students could be expected to know) in order to make sense of the lesson. And, as described immediately below, the teacher's agenda for this lesson was much more open than were the agendas in any of the others.[26]

Ball's class had been studying the properties of even and odd numbers. Among the things that had taken place:

* The class had arrived at a working definition of what it means for a number to be even. In one student's words, "If you have a number that you can split up evenly without . . . having to split one in half, then . . . it's an even number." Thus, for example, 8 (which equals 4 plus 4) is even. In contrast, 5 can not be split into two equal piles of whole numbers, so 5 is not even.

* Some members of the class had noted the alternation of even and odd numbers on the number line.

* The previous day there had been a special event, in which Ball's third graders and some fourth graders met together to discuss the properties of even and odd numbers. One issue that came up at that meeting was whether the number zero was even, odd, or "special."

---

[26]Note that our comparison is of the lessons and not of the teachers. At different times of the year, Schoenfeld and Minstrell might have taught lessons with many of the characteristics of the lesson taught by Ball, and vice-versa. Just as (say) Minstrell's teaching is likely to look very different in a review lesson and an exploratory lesson, the models of his teaching in those two contexts are likely to look very different. In every case, we are modeling a teacher engaged in a particular type of lesson, with particular goals, knowledge base, etc.

* One student had noted the previous day that some even numbers were "made up of" even numbers - e.g., that is 8 is 4 plus 4. It was not clear at the time whether he meant that adding an even number to itself gives an even number, that some but not all even numbers yield an even number when split in two, or that all even numbers yield an even number when split in two. The issue had not been pursued.

As will be clear from her opening comments, one of Ball's primary goals that day was to have students reflect on the previous day's activities - not just on the mathematics but, more importantly at the beginning of the lesson, on the state of their knowledge and learning. Among Ball's overarching goals for herself and for her students was the development of a particular kind of intellectual community in which the pursuit of mathematical ideas in reasoned ways was accorded great value. Ball herself, as a teacher, had as a goal understanding as much as she could about each student's understanding of the mathematics that they were studying. As seen below, the presence of such goals can lead to dilemmas and tensions while teaching. Suppose, for example, that a student makes a comment that raises an important content issue, points to an unresolved issue from the day before, and raises issues about the way the student is thinking about the content. Which of these paths might be pursued, to what advantage and with what costs? Juggling these concerns is a delicate act.

Ball began the lesson with the following comments:

I'd like to open, open the discussion today with um - I have a few questions about the meeting yesterday that I'd like to ask. So, to begin with, I would just like everybody to put pens down, there's nothing to take notes about or do right now. But I'd like you to be thinking back to yesterday and to the meeting that we had on even and odd numbers and zero. And I have a few questions. First - my first question is, I'd just like to hear some comments about what you thought about the meeting, what you noticed about the meeting, what you learned at the meeting, just what kinds of comments you have about yesterday's meeting? And could you listen to one another's comments, so that we can um, benefit from what other people say? See what you - what you think about other people's comments? [Ball names a particular student], do you want to start?

Ball's intentions are clear from what she said. She followed up on these intentions, as exemplified by the following two snippets of dialogue:

Ball: Was there an example of something yesterday that you understood a little bit more during the meeting?

Student: Well, I didn't think that zero was - zero, um - even or odd until yesterday they said that it could be even because of the ones on each

side is odd, so that couldn't be odd. So that helped me understand it.

Ball: Hmm. So you thought about something that came up in the meeting that you hadn't thought about before? Okay.

and

Student: I thought that zero was always going to be a even number, but from the meeting I sort of got mixed up because I heard other ideas I agree with and now I don't know which one I should agree with.

Ball: Um-hm. So what are you going to do about that?

Student: Um, I'm going to listen more to the discussion and find out.

But, as students commented on the meeting some of their comments touched on their understandings. Ball found herself pursuing unresolved issues related to those understandings:

Student: Um, first I said that um, zero was even but then I guess I revised so that zero, I think, is special because um, I - um, even numbers, like they *make* even numbers; like 2, um, 2 makes 4, and 4 is an even number; and 4 makes 4; 4 is an even number; and um, like that. And, and go on like that and like 1 plus 1 and go on adding the same numbers with the same numbers. And so I, I think zero's special.

Ball: Can I ask you a question about what you just said? And then I'll ask people for more comments about the meeting. Were you saying that when you put even numbers together, you get another even number -

Student: Yeah.

Ball: - or were you saying that all even numbers are made up of even numbers?

Ball pursued this issue until it was resolved and then returned to the topic of her announced goal:

Ball: More comments about the meeting? I'd really like to hear from as many people as possible what comments you had or reactions you had to being in that meeting yesterday. Shea?

Shea: Um, I don't have anything about the meeting yesterday, but I was just thinking about 6, that it's a . . . I'm just thinking. I'm just thinking it can be an odd number, too, 'cause there could be 2, 4, 6, and 2, three 2s, that'd make 6 . . .

After a brief interaction, in which she thought she had clarified the issue for Shea, Ball opened things once again: "Other people's comments?" The student who next commented focused back on Shea's statement, arguing that 6 cannot be an odd number because of the pattern "even, odd, even, odd, . . ." on the number line. After another comment on the same issue, Ball found herself focusing on mathematical content as she tried to sort things out with Shea: "What's our working definition of an even number? Do you remember from the other day the working definition we're using?" The class pursued that definition for a while, ultimately agreeing that 6 satisfied the working definition. Then the following exchange took place:

Ball: One of the points here is that if it fits the definition then we would call it even. If it fits our working definition, then we would call it even.

If it can fit the um, definition for odd, too?

Shea: What is the definition for odd, maybe we need to talk about that.

Ball: There could be um a - can I give it? Um, there could be like um, 6 - there

Shea: could be like, there's only three, three numbers to make an odd number in there and not two numbers . . .

Is that what other people are assuming about the definition of an odd

Ball: number? Very interesting. We didn't discuss what our definition was for odd numbers.

Later it became clear that Ball and most of the students had assumed that if a number is even it cannot be odd, and vice-versa - but that Shea had not made that assumption. It is, in fact, quite reasonable mathematically that if one has a non-standard definition of what it means to be odd, some numbers (6 in particular) might satisfy the definition of being both an even and an odd number; note that some numbers are both squares and cubes, for example.[27] In any case, the class was then engaged in a discussion of mathematical content. What transpired as the class went on to pursue various definitions of oddness and their entailments is fascinating.

To this point, the Teacher Model Group has had a difficult time producing a comprehensive analysis of the segment of the lesson just described. The issue is not that it's hard to produce a parsing at the phenomenological level, or that it's difficult to tell a good narrative story about what happened - doing each of those is straightforward. Nor is it difficult to attribute beliefs to Ball - she, like Minstrell, has written extensively about her work, and has made available a corpus of data that can be used to triangulate on the attribution of beliefs. (She

[27]Shea's definition of "even" was the class's working definition of even, that an even number is one made up of *two* of the same things. His claim was that since 6 is also made up of *three* of the same things (*three* 2s), and three is odd, that 6 is odd. Ultimately, his proposed definition was that a number is odd if it is made up of an odd number of the same things.

has also volunteered to discuss these issues with us at an appropriate time.) The difficulty is with the attribution and prioritization of goals and action plans in a way that produces a coherent account of why Ball did what she did - and not simply an ad hoc justification of it. As noted above, Ball has multiple and often competing goals in this lesson, and she balances among them. We have yet to produce what we consider to be a consistent and compelling account of why things evolved in the particular ways they did. We may yet do so, for there is a substantial amount of additional data available to us.[28] But it will not be easy. This particular lesson has lived up to our expectations in testing our capacity to model it. It raises serious issues concerning how well we can model lessons that have multiple and sometimes conflicting goals, and agendas that evolve radically in response to what students say and do.

[28]For methodological reasons, we have been parsimonious in the data we have allowed ourselves to use. The idea was to start with a small amount of data - the tape and transcript of the class, and some other background information - and see if there were claims we felt we could make on the basis of those data. If those claims turned out to be solid (i.e., confirmed by Ball and consistent with the additional data), that would be evidence in favor of the methods we were using. If, however, we made claims that turned out to be wrong, we would then have grounds for concern regarding our analytic methods. That is how things worked (with positive results) during the analysis of the Minstrell benchmark elaboration lesson. Here, the nature of the lesson is such that we do not feel confident about making strong claims given the data we started with. In consequence, we will soon begin to look at additional class tapes and other related materials.

8. Discussion: Strengths, Limits, and Implications; Next steps; and, Does this enterprise make sense?

I begin with a brief discussion of the status of and prospects for the theory and the model. Following that discussion I turn to the bottom line question - even if the enterprise can succeed, why should anybody care?

8.1. How good are the theory and the model?

Let me begin with a reminder of the enterprise as a whole. As described in the introduction, "Our intention is to provide a detailed theoretical account of how and why teachers do what they do 'on line' - that is, while they are engaged in the act of teaching. This theoretical characterization of teaching is embodied in a model of the teaching process. The model describes, at a level of mechanism, the ways in which the teacher's goals, beliefs, and knowledge interact, resulting in the teacher's moment-to-moment decision-making and actions."

To assess the status and prospects of this enterprise I return to the three criteria introduced in Section 2 of this paper: explanatory power, predictive power, and scope.

Explanatory power

Perhaps the greatest strengths of the work lie along this dimension. If there has been one major contribution of cognitive science over the past two decades, it has been its focus on the explanations at the level of mechanism - descriptions that say how things fit together and how they work. This work lives squarely within that tradition. Make no mistake, the analyses require a huge amount of work. But the payoff is a fine-grained description not only of what happens, but why.

As one example, consider the analysis given in Sections 5.4 and 5.5 of the part of Mark Nelson's lesson in which his class discussed reducing the rational expression $\frac{x^5}{x^3}$. At the most crude descriptive level, the story of what took place is simple: Nelson tried something that didn't work, got stuck, perseverated, and then made an exit when he was able to. The detailed blow-by-blow descriptive version is given in Section 5.3. A typical cognitive analysis would push things further, suggesting why he got stuck. It would note that (a) when Nelson made a slash mark through two xs at the board, "canceling" them, making that slash mark represented the division of an x by an x; the result being 1; and (b) the meaning of that canceling was not made explicit to his students, and he lost them as a result. But this analysis goes further. It shows how his epistemological stance (which centers on teaching as explanation rather than as the exploration of student understandings), his goals (including student engagement), his lesson image (which involved using students' answers to his questions as springboards for explanations), and his knowledge base (strong in some ways but lacking in

specific knowledge of the difficulties his students were likely to encounter or ways to deal with those difficulties) interacted with each other. It describes those interactions at a very fine grain size that serves to explain what he did throughout the lesson and why, at a level of mechanism. Indeed, the model actually predicts the bind Nelson finds himself in, and why it is so difficult for him to extricate himself from it.

The analysis of Minstrell's lesson in Section 6 (and in the more extended study it draws from, Schoenfeld, Minstrell, & van Zee, 1996) provides another, albeit very different, case in point. I will not repeat the discussion, which was summarized in Section 6.5. But the bottom line is that the analysis of Minstrell's beliefs, goals, and knowledge explains and ofttimes predicts his actions at a line-by-line level of interactions. It does so not only when things are going as planned, but (cf. the discussion of [1.2.2.3] in Section 6.4) when unexpected events take place. In short, the model lives up to its promises (in these two cases at least): it explains in fine-grained detail how and why both teachers did what they did, on the fly, in the classroom.

Predictive power

One must begin with the necessary caveat: Human behavior, especially in complex social settings, is not predictable in the sense of predictions made in the physical sciences. However, it is possible to make predictions of a certain type using the kind of model developed here, and to judge the model by the accuracy of those predictions.

The cases discussed in this paper provide a number of examples that are essentially predictive in nature. I highlight that aspect of those examples briefly in this discussion. First, the analysis of Nelson's belief system and knowledge base (specifically, his pedagogical strategies and pedagogical content knowledge) suggests that he will be oriented toward providing explanations in a particular way and that he will get stuck when the students fail to "get" his explanation. Second, the analysis of Minstrell's beliefs allows for the firm prediction that he will pursue the "complicated" suggestion made by a student - and it suggests how he will do so. Third, the decision procedures represented in Figures 6 and 10 provide strong bases for making predictions about what Minstrell and I, respectively, will do when engaging students in discussions of particular problems. In each of these cases, you could run a videotape up to a point where something happens (e.g., when students' answers to a question Nelson poses do or do not provide him with particular kinds of information, or a student in Minstrell's class makes a non-standard but relevant suggestion). Then you could run the model up to that point, and ask what the model would do next. Depending on the context, you may not get just one answer, of course - but the model will indicate what the teacher is very likely to do, and what the teacher is very unlikely to do. You can then compare the model's predictions with what actually happened. (I note that for the lessons analyzed in this paper, the model

does quite well. How it will fare on the analysis of the "Shea number" lesson remains to be seen.)

Of course, predictions based on tapes that have already been analyzed are the easiest kind to make - though a parsimonious model that covers a lot of territory, such as the one in Figure 6, is still non-trivial. Nonetheless, one might look to make more "distant" predictions. To do so one must recall that the model is highly context-specific. For example, both Minstrell's lesson and mine took place at the beginning of the year, with specific goals in mind; the degree of teacher- directedness in those lessons is much stronger than later in the year, when the classroom community is well established and functioning on its own steam. Since the teachers would not act the same across the boards, one would not expect to use a model of the teacher early in the year to predict behavior in mid year. But, there are predictions one could make. For example, one might well be able to make predictions about Minstrell or me teaching the opening days of our respective courses in a different year. Or, suppose that someone has used the model to analyze a sequence of a few days' lessons in a course. Presumably that person would be in a good position to make predictions about events in the next day of the course. (That is, he or she analyzes days k through n, and watches a video of day $(n + 1)$ for the first time, up to a point where something interesting happens. He or she then makes a prediction about what happens next.) Finally, there are some aspects of a teacher's beliefs, values, and style that are so consistent that they allow for predictions even in lesson contexts that are unfamiliar. For example, the following is a reasonably safe bet at almost any point in Minstrell's course. Suppose a student raises an unexpected issue whose exploration would require a detour from Minstrell's lesson plan. Suppose in addition that Minstrell perceives the issue as being relevant to what is being discussed and that exploring it would provide an opportunity for the class to engage in sense-making. Then Minstrell will, constraints permitting, make time for the detour.

I don't want to make too big a deal of prediction. The goal is not to be scientistic here, but to defend against producing ad hoc (post hoc) explanations. If one's analyses allow for the plausible foretelling of events to come, then one has some reason to be confident of the analyses. Admittedly, the lessons discussed in this paper are of a somewhat limited range: Nelson ran out of resources early, and Schoenfeld's and Minstrell's lessons are ones in which the teacher had a very elaborate lesson image and the means by which to carry it out. Nonetheless, in those circumstances, the model does reasonably well on a predictive basis.

Scope

Scope is the criterion about which I have my biggest questions. The questions come in two different dimensions, internal and external to the model.

Issues internal to the model.

As described in the body of this paper, the Teacher Model group has made a deliberate choice of test cases to see how far the model will stretch. The first analysis focused on a traditional teacher-centered lesson taught by a relatively new teacher. In contrast, the second analysis focused on a decidedly non- traditional student-centered lesson taught by a highly experienced teacher. These differences were not the least bit problematic. We were able to model both cases without difficulty, and the comparison of what happened when both teachers ran into unexpected circumstances was illuminating. The third case, the analysis of my problem solving instruction, was initially independent of the Teacher Model work, and proceeded in parallel with it. Because there are some similarities between Minstrell's initial lesson and mine, the third case did not push the boundaries of the space defined by the first two - but it provided a nice confirmation of much of the Minstrell analysis. (Moreover, the confirmation is at the deep rather than the surface level. Minstrell's style and mine are very different, and parallels such as the similarity between Figures 6 and 10 are anything but obvious in a side-by-side comparison of the two chunks of instruction.) With the fourth case, Deborah Ball's lesson, we returned to the deliberate strategy of pushing the model (and ourselves) as far as we could.

As the discussion in Section 7.2 indicates, modeling that lesson has been a significant challenge. That is the case for at least two reasons. The first is that the class takes place mid-year, and there is a lot more "history" between teacher and students to take into account than in our other analyses. This could turn out to be a significant problem, but my best guess is that it will not. We have been parsimonious in seeking data, and I suspect that when we look at more data the issues that depend on our knowing what took place in previous days and weeks will be resolved. The second and more problematic issue has to do with the evolving agenda for the class. There is no doubt that in this class Ball plays a major role in shaping the conversations and the directions they take - but she is also tremendously responsive to issues raised by the students and is willing to interrupt some agenda items to pursue, temporarily, other issues of interest. The juggling of priorities on her part makes it a challenge for us to assign, in a consistent way, levels of activation to constellations of beliefs, goals, and knowledge that influence her actions. In consequence, parsing the lesson into chunks and providing a goal and action plan analysis that corresponds to those chunks has been much more difficult than in the three previous cases. There is some reason to believe that with more information, we can get this sorted out. Below I briefly describe the implications if we can, and if we can not.

There are some interesting distinctions and similarities between the Nelson and Minstrell lessons. Nelson's lesson can be described as teacher-centered, while Minstrell's is better described as student-centered - but both lessons have the property that the teacher sets the agenda and that the action follows the teacher's

lead.[29] No matter what the outcome of the effort to analyze Ball's tape, it is clear that the model can deal very well with lessons of the type just described. (Note: There is a larger data base to ground this assertion than just the lessons discussed in this paper. We have, of course, looked at a fair number of lessons other than the ones described here - although often for different purposes and not at the level of detail necessary to model them. Tapes we have examined included lessons taught by experienced traditional teachers dealing with familiar and traditional curricular materials, and also with materials that were innovative and new to them. Those teachers, too, set the classroom agenda and had the class follow their lead. Though we did not perform detailed analyses of those lessons, it was clear that doing so would merely be a matter of time and effort.)

In sum, there is reason to be confident that we can model classroom sessions with clearly teacher-led agendas - at least, classrooms where the lesson plays out somewhat in accord with the teacher's lesson image and/or shifting agenda. That is certainly not all lessons, but it's not an inconsequential percentage of them. That's the worst case for the model's scope, and that's not bad. If we succeed in modeling the Shea number lesson, which is of a very different type, then the scope of the model is likely to be significantly larger - though how much of the whole territory we cover is still open to question. While the agenda shifts quite a bit during the Shea number lesson, Ball is still a powerful force in shaping the directions the discussion takes. We may or may not be able to model what takes place in classrooms where the agenda is co-constructed interactively by teacher and students. Moreover, as the locus of decision-making shifts to include the students more, the centrality of the model to explain "what happened" in the classroom may diminish significantly. But that is an issue for the next subsection, dealing with external issues.

Issues external to the model.

As noted, ours is an ambitious and yet constrained research agenda. The ambition is clear and has been elaborated at length. But now it is time to return to the constraints. This paper describes a theory of teaching-in-context. The notion is that with a complete enough analysis of a teacher's beliefs, goals, and knowledge base, we can describe/predict how the teacher will act in the classroom. This kind of approach entails two very strong constraints.

The first constraint is the model's very strong focus on the teacher. The following is obvious, but it needs to be said: an awful lot happens in classrooms above and beyond the thoughts and actions of teachers. The model discussed here does not address classroom reality from the students' point of view, for example; nor does it address, except through the teacher's eyes, issues such as the formation of the classroom community and its growth and change. The field needs to address

[29]Once again, it should be stressed that I am talking about specific lessons and not about general traits of the teachers. The teachers might well behave very differently in different circumstances.

such issues. Modeling individual students is well within the scope of current work, as is modeling the teacher, but we lack the theoretical frames with which we can provide powerful and coherent descriptions of a community and its growth, intertwining the stories of the individuals and the collective.

The second constraint is the fact that the model limits its attention to what takes place in the classroom. The model embodies a theory of action-in-context, providing snapshots of behavior in particular contexts over relatively short time spans. It says that if we know the teacher's beliefs, goals, and knowledge base, we can understand how the teacher is likely to act *now*. Imposing that constraint makes the current task - understanding and modeling what the teacher does and why, in fine-grained detail - manageable although still challenging. But we should not forget just how strong a constraint that is. What takes place inside a classroom is shaped by myriad forces outside it. The kinds of sociological, anthropological, and historical approaches that shed so much light on those forces tend to operate at different levels of grain size, with different standards of explanation and, at times, different underlying epistemological stances than the ones that are at the core of the work described here. To make further progress we must find ways to reconcile some of those differences, or develop perspectives and methods that extend far beyond the perspectives and methods that are currently available.

8.2. Why should anyone care about building this kind of model?

The model described in this paper is complex, and it takes a lot of work to model any particular instructional segment. Why should anyone bother to do the kind of detailed modeling exercise discussed here?

There are, I think, a number of important reasons.

One main reason is theoretical: models of this type represent an advance in understanding human cognition. Having been at this game for a quarter century or so, I can marvel at the progress we have made, and can see this kind of work as representing a "next step" toward the grand goal of understanding human behavior in complex social settings. It has been a phenomenal journey. In the 1970s, we were just beginning to make sense of problem solving behavior. It was all we could do to analyze short problem solving sessions that took place in laboratory settings, where individuals worked in isolation. The focus of research was on skills and strategies, and what we now consider to be fundamental concepts such as metacognition and belief systems were barely on the intellectual horizon. Back then the theoretical constructs available to the research community had severely limited explanatory power. Our methodological tool kit was sparse and our capacity to think about thinking was so constrained that we studied people in isolation because we were unable to deal with the dynamics of social interactions. What progress we have made! As the discussion of "scope" above indicates, we still have a long way to go. But, we are making good

progress toward understanding how and why people do what they do, where it counts - in interactions with others.

Another main reason, which may sound theoretical but is ultimately deeply pragmatic, concerns theories of competence and their potential impact on practice. The general argument is simple: if you want to help people do a job well, you have to know what it takes to do the job. I think a good case can be made that if you really want to understand teaching well enough to make a difference, you will have to do something akin to what we have been doing.

Let me briefly reprise the argument made in Section 1.3. Consider the long-term impact of research on mathematical thinking and learning. It took decades of research, but ultimately our evolving conception of "what counts" in understanding and doing mathematics has made a big difference. The current conception of what it means to think mathematically differs radically from that of thirty years ago. Back then, the focus was on content "mastery." (Indeed, during the heyday of the behaviorists the very idea of mental processes - the foundation of most current work - was considered theoretically untenable.) Today's definition of thinking mathematically includes having a solid knowledge base, but it also includes knowing a wide range of problem solving strategies, having modeling skills, metacognitive skills, productive beliefs, and more. This reformulation of mathematical competence has led to the redesign of mathematics curricula: over the past half dozen years the National Science Foundation has funded a series of curriculum projects whose goal it is to help students develop a solid knowledge base, facility with problem solving strategies, good modeling and metacognitive skills, and so on. Those curricula are now making their way into the marketplace. In short, theoretical advances are at the point of having a profound impact on practice (after some 25 years).

In a similar way, having a deeper understanding of teaching should have real payoffs in the long run. But even if we grant that point, one should ask if it is necessary to build the kinds of models that are the focus of this paper. What does this kind of model offer that we didn't already know? We know, for example, that teachers' beliefs are important (see, e.g., Thompson, 1992). We know that professed beliefs don't tell the whole story - the story of Mrs. Oublier (Cohen, 1990) makes that clear. We know that what teachers envision (see, e.g., Clandinin, 1986, Morine-Dershimer, 1978-79) and what they know (Carter & Doyle, 1987; Calderhead, 1996; Leinhardt, 1993) make a difference. We know that teachers juggle multiple and competing demands (Lampert, 1985), and can be limited by the resources they are capable of bringing to bear in the classroom (Borko, Eisenhart, Brown, Underhill, Jones, & Agard, 1992). What is the value added of this kind of model?

One answer is that it is not enough to know how beliefs, envisionings, and the knowledge base work separately; we have to know how they interact. We know, for example, that beliefs are important. But, two teachers may profess the same

beliefs and act very differently in similar situations. Or, two teachers may appear to have the same goals - even a shared lesson plan - and yet the lessons will play out very differently. Why? Because different beliefs may result in the same goals being interpreted in different ways, and the ways in which goals actually play out (via the selection of action plans) will depend crucially on the nature of each teacher's knowledge base.

A second and related answer concerns the practical entailments of the first. As the model indicates, beliefs, goals, and resources are deeply intertwined; meaningful change will involve a realignment of all of these. Many short-term professional development efforts fail because they do not address all the components necessary for change - they may provide motivation but not the chance for teachers to develop the necessary cognitive resources, or might offer skills or materials that will be used very differently than intended because of the teacher's belief structure. It should be clear that developing a better understanding of how beliefs, goals, and knowledge interact will help us to get a better sense of what it takes to support teacher change.

A third answer is that explanation at the level of mechanism really does make a difference. The fine-grained analysis of Nelson's lesson segment does much more than tell us that he ran into difficulties. The analysis tells us why he got into trouble and why he was unable to extricate himself from it, and it suggests what kinds of perspectives and competencies he would need to develop in order to avoid such difficulties in the future. Likewise, the analysis of Minstrell's lesson segment does more than tell us that he was successful at achieving certain goals. It enables us to understand why he was able to do what he did, specifying the support structure that enabled him to interact with his students in a manner consistent with his beliefs and goals. If you hope to help teachers to do something similar, understanding "what it takes" is a good place to start. (It is worth noting that although the analyses are typically complex, their end products need not be. Figures 6 and 10 are cases in point, in that they provide relatively simple structural descriptions of very complex behavior. In applied terms, one could imagine using something like Figure 6 as a heuristic guide for developing the kinds of questioning skills evidenced by Minstrell.)

A fourth answer is more prospective but ultimately more important. The model allows us to take good "snapshots" of people's teaching, in that it provides a rich and detailed portrayal of what they do in the classroom and why. One thing we lack, and we need desperately, is a sense of teachers' developmental trajectories. What skills, understandings, and beliefs develop in what ways, and how do they interact? Are there general developmental patterns that we can understand and perhaps influence through pre- and in-service professional development? Are there developmental plateaus that might be avoided, or points in teachers' development where specific kinds of interventions might be particularly useful? The more we know about teachers' evolution as professionals, the more we can develop ways to shape that evolution productively.

Are these projections of improved practice mere pie-in-the-sky optimism? I don't think so, although the time and effort required to make real progress should not be underestimated. As noted above, it took a good 25 years to reconceptualize what it means to think mathematically, and to develop instructional materials along the lines suggested by that reconceptualization. But we're now at that point, reaping the rewards of the basic research. In contrast, we are still at the very beginning of the learning curve with regard to our understanding of teaching. That means we have a very interesting and, I believe, ultimately a very productive journey ahead of us.

Acknowledgment

I would like to express my thanks to the members of the Teacher Model Group at Berkeley for the pleasures of collaboration that led to this paper and for their many contributions to it. Members of the group during the time this paper was produced included Julia Aguirre, Tu Banh, Gina Banzon, Jake Disston, Lisa Eagan, Ilana Horn, Andrew Izsak, Cathy Kessel, Brett Lorie, Rodrigo Madanes, Liping Ma, Susan Magidson, Jim Minstrell, Alexis Nogelo, Manya Raman, Eric Saterlie, Alex Sheffler, Miriam Gamoran Sherin, Natasha Speer, Emily van Zee, Joseph Wagner, Eunice Yi, and Daniel Zimmerlin. Special thanks go to Cathy Kessel, Daniel Zimmerlin, Emily van Zee, Gaea Leinhardt, Julia Aguirre, Ilana Horn, Miriam Gamoran Sherin, Rich Shavelson, and Natasha Speer.

REFERENCES

Aguirre, Julia, & Speer, Natasha. (1996). Examining the relationship between beliefs and goals in teacher practice. Paper presented at the annual meeting of the American Educational Research Association, New York, April 8-12, 1996.

Arcavi, Abraham, Kessel, Cathy, Meira, Luciano, & Smith, John P. (1998). Teaching mathematical problem solving: An analysis of an emergent classroom community. In Alan H. Schoenfeld, James J. Kaput, & Ed Dubinsky (Eds.),*Research in Collegiate Mathematics Education, Volume III* (pp. 1-70). Washington, DC: Conference Board of the Mathematical Sciences.

Ball, Deborah. (1988). *Knowledge and reasoning in mathematical pedagogy: Examining what prospective teachers bring to teacher education.* Unpublished doctoral dissertation, Michigan State University.

Ball, Deborah, & Lampert, Magdalene. (in press). Multiples of evidence, time, perspective: Revising the study of teaching and learning. To appear in: Lee S. Shulman & Ellen C. Lagemann (Eds.), *Commission on the Improvement of Educational Research: Papers on strategies, training, and communication (working title).* Stanford, CA: National Academy of Education.

Berliner, David C. (1994). Expertise: The wonder of exemplary performances. In John M. Mangier & Cathy Collins Block (Eds.), *Creating powerful thinking in teachers and students: Diverse perspectives* (pp. 161-186). Fort Worth, TX: Holt, Reinhart, & Winston.

Berliner, David C., & Calfee, Robert C. (Eds.). (1996). *Handbook of educational psychology* . New York: Macmillan.

Borko, Hilda, Eisenhart, Margaret, Brown, Catherine, Underhill, Robert, Jones, Doug, & Agard, Patricia. (1992). Learning to teach hard mathematics: Do novice teachers and their instructors give up too easily? *Journal for Research in Mathematics Education* , 23(3), 194-222.

Borko, Hilda, & Putnam, Ralph. (1996). Learning to teach. In David C. Berliner & Robert C. Calfee (Eds.), *Handbook of educational psychology* (pp. 673-708). New York: Macmillan.

Bruner, Jerome. (1986). *Actual minds, possible worlds.* Cambridge, MA: Harvard University Press.

Calderhead, James. (1988). The development of knowledge structures in learning to teach. In James Calderhead (Ed.), *Teachers' professional learning* (pp. 51-64). London: Falmer.

Calderhead, James. (1996). Teachers: Beliefs and knowledge. In David C. Berliner & Robert C. Calfee (Eds.), *Handbook of educational psychology* (pp. 709-725). New York: Macmillan.

Carter, Kathy, & Doyle, Walter. (1987). Teachers' knowledge structures and comprehension processes. In James Calderhead (Ed.), *Exploring teachers' thinking* (pp. 147-160). London: Cassell.

Cazden, Courtney. (1988). *Classroom discourse.* Portsmouth, NH: Heinemann.

Clandinin, D. Jean. (1986). *Classroom practice: Teacher images in action.* Philadelphia: Falmer Press.

Clark, Christopher, & Peterson, Penelope. (1986). Teachers' thought processes. In Merlin C. Wittrock (Ed.), *Handbook of research on teaching* (Third edition, pp. 255-314). New York: Macmillan.

Clark, Christopher, & Yinger, Robert. (1987). Teacher planning. In J. Calderhead (Ed.), *Exploring teachers' thinking* (pp. 84-103). London: Cassell.

Cohen, David. (1990). A revolution in one classroom: The case of Mrs. Oublier. *Education Evaluation and Policy Analysis*, 12(3), 311-329.

Connelly, F. Michael, & Clandinin, D. Jean. (1985). Personal practical knowledge and the modes of knowing: Relevance for teaching and learning. In Elliot Eisner (Ed.), *Learning and teaching the ways of knowing* (pp. 174-198). Chicago: University of Chicago Press.

Cooney, Thomas. (1985). A beginning teacher's view of problem solving. *Journal for Research in Mathematics Education*, 16(5), 324-336.

Darling-Hammond, Linda. (Ed.). *Review of research in education*, Volume 20. Washington, DC: AERA.

diSessa, Andrea A. (1993). Toward an epistemology of physics. *Cognition and Instruction, 10* (2-3), 105-225.

Elbaz, Freema. (1983). *Teacher thinking: A study of practical knowledge.* New York: Nichols.

Ernest, Paul. (1989). The knowledge, beliefs, and attitudes of the mathematics teacher: A model. *Journal of Education for Teaching*, 15(1), 13-33.

Fennema, Elizabeth, & Franke, Megan. (1992). Teachers' knowledge and its impact. In D. Grouws (Ed.), *Handbook of research on mathematics teaching and learning* (pp. 147-164). New York: Macmillan.

Fenstermacher, Gary. (1994). The knower and the known: The nature of knowledge in research on teaching. In Linda Darling-Hammond (Ed.), *Review of research in education*, Volume 20 (pp. 3-56). Washington, DC: AERA.

Frederiksen, John, Sipusic, Michael, Gamoran, Miriam, & Wolfe, Ed. (1992). *Video portfolio assessment*. Emeryville, CA: Educational Testing Service.

Leinhardt, Gaea, & Greeno, James G. (1986). The cognitive skill of teaching. *Journal of Educational Psychology, 78* (2), 75-95. Also appears in P. Goodyear (Ed.) (1991), *Teaching knowledge and intelligent tutoring*. Norwood, NJ: Ablex.

Grouws, Douglas A. (Ed.). (1992). *The handbook of research on the teaching and learning of mathematics*. New York: MacMillan.

Grossman, Pamela L, & Stodolsky, Susan S. (1994). Considerations of content and the circumstances of secondary school teaching. In Linda Darling-Hammond (Ed.), *Review of Research in Education*, Volume 20 (pp. 179-222). Washington, DC: AERA.

Hinsley, Dan A., Hayes, John R. & Simon, Herbert A. (1977). From words to equations: Meaning and representation in algebra. In P. A. Carpenter & M. A. Just (Eds.), *Cognitive processes in comprehension* (pp. 89-106. Hillsdale, NJ: Erlbaum.

Lampert, Magdalene. (1985). How do teachers manage to teach? Perspectives on problems in practice. *Harvard Educational Review, 55* (2), 178-194.

Lampert, Magdalene. (1990). When the problem is not the question and the solution is not the answer: Mathematical knowing and teaching. *American Educational Research Journal,* 27(1) 29-63.

Leinhardt, Gaea. (1988). Situated knowledge and expertise in teaching. In James Calderhead (Ed.), *Teachers' professional learning* (pp. 146-168). London: Falmer.

Leinhardt, Gaea. (1993) On teaching. In Robert Glaser (Ed.), *Advances in Instructional psychology* (Volume 4, pp. 1-54). Hillsdale, NJ: Erlbaum.

Leinhardt, Gaea, & Greeno, James G. (1986). The cognitive skill of teaching. *Journal of Educational Psychology, 78* (2), 75-95. Also appears in P. Goodyear (Ed.) (1991), *Teaching knowledge and intelligent tutoring* . Norwood, NJ: Ablex.

Leinhardt, Gaea, Weidman, C., & Hammond, K. M. (1987). Introduction and integration of classroom routines by expert teachers. *Curriculum Inquiry, 17* (2), 135-176.

Lewis, Charles. (1974). *The juggling book, by Carlo* . New York: Vintage Books.

Ma, Liping. (1996). *Profound understanding of fundamental mathematics: what it is, how it is attained, and why it is important.* Unpublished doctoral dissertation, Stanford University.

Mehan, Hugh. (1979). *Learning lessons: Social organization in the classroom.* Cambridge, MA: Harvard University Press.

Minstrell, James A. (1982). Explaining the "at rest" condition of an object. *The Physics Teacher* 20, 10-14.

Minstrell, James A. (1989). Teaching science for understanding. In: Lauren B. Resnick & Leopold E. Klopfer (Eds.), *Toward the thinking curriculum: Current cognitive research* (pp. 129-149). Washington, DC: Association for Supervision and Curriculum Development.

Minstrell, J. (1992). Facets of students' knowledge and relevant instruction. In R. Duit, F. Goldberg, & H. Niedderer (Eds.), *Research in physics learning: Theoretical issues and empirical studies* (pp. 110-128). Kiel, Germany: Institute for Science Education.

Minstrell, Jim, & Stimpson, Virginia. (1996). A classroom environment for learning: Guiding students' reconstruction of understanding and reasoning. In Leona Schauble & Robert Glaser (Eds.), *Innovations in learning: New environments for education* (pp. 175-202). Mahwah, NJ: Erlbaum.

Morine-Dershimer, Greta. (1978-79). Planning in classroom reality: An in-depth look. *Educational Research Quarterly, 3* (4), 83-89.

Pajares, M. Frank. (1992). Teachers' beliefs and educational research: Cleaning up a messy construct. *Review of Educational Research, 62* (3), 307-332.

Pólya, George. (1945). *How to solve it.* Princeton: Princeton University Press.

Schank, Roger, & Abelson, R. (1977). *Scripts, plans, goals, and understanding.* Hillsdale, NJ: Erlbaum.

Schoenfeld, Alan H. (1985). *Mathematical problem solving* . Orlando, FL: Academic Press.

Schoenfeld, Alan H. (1988). When good teaching leads to bad results: the disasters of "well taught" mathematics classes. *Educational Psychologist,* 23(2), 145-166.

Schoenfeld, Alan H. (1998). Reflections on a course in mathematical problem solving. In Alan H. Schoenfeld, James J. Kaput, & Ed Dubinsky (Eds.),*Research in Collegiate Mathematics Education, Volume III* (pp. 81-113). Washington, DC: Conference Board of the Mathematical Sciences.

Schoenfeld, Alan H. (1996). Elements of a model of teaching. Paper presented at the annual meeting of the American Educational Research Association, New York, April 8-12, 1996.

Schoenfeld, Alan H., Minstrell, Jim, & van Zee, Emily (1996). The detailed analysis of an established teacher carrying out a non-traditional lesson. Paper presented at the annual meeting of the American Educational Research Association, New York, April 8-12, 1996.

Schoenfeld, Alan H., Smith, John P., & Arcavi, Abraham A. Learning. (1993). In R. Glaser (Ed.), *Advances in Instructional Psychology* (Volume 4, pp. 55-175). Hillsdale, NJ: Erlbaum.

Schön, Donald A. (1982). *The reflective practitioner.* New York: Basic Books.

Shavelson, Richard. (1986). Interactive decisionmaking: Some thoughts on teacher cognition. Paper presented at I congreso internacional, "Pensamientos de los profesores y toma de decisiones," Seville, Spain, June 19-21, 1986.

Sherin, Miriam G., Sherin, Bruce, & Madanes, Rodrigo. (1996). Where terms collide: Clarifying diverse accounts of teacher knowledge. Paper presented at the annual meeting of the American Educational Research Association, New York, April 8-12, 1996.

Shulman, Lee S. (1986). Those who understand: Knowledge growth in teaching. *Educational Researcher, 17* (1), 4-14.

Shulman, Lee S. (1987). Knowledge and teaching: Foundations of the new reform. *Harvard Educational Review,* 57, 1-22.

Stigler, James, Fernandez, Clea, & Yoshida, Makoto. (1996). Cultures of mathematics instruction in Japanese and American elementary classrooms. In Thomas P. Rohlen & Gerald K. LeTendre (Eds.), *Teaching and learning in Japan* (pp. 213-247). Cambridge: Cambridge University Press.

Stodolsky, Susan S. (1985). Telling math: Origins of math aversion and anxiety. *Educational Psychologist, 20* (3), 125-133.

Strauss, Sidney, & Shilony, Tamar. (1994). Teachers' models of children's minds and learning. In Lawrence A. Hirschfeld and Susan A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture* . Cambridge: Cambridge University Press.

Talbert, Joan E., McLaughlin, Milbrey W., & Rowan, Brian. (1993). Understanding context effects on secondary school teaching. *Teachers College Record* , 95(1), 45-68.

Thompson, Alba. (1992). Teachers' beliefs and conceptions: A synthesis of the research. In D. Grouws (Ed.), *Handbook of research on mathematics teaching and learning* (pp. 127-146). New York: Macmillan.

van Zee, Emily H., & Minstrell, James A. (1997a). Reflective discourse: Developing shared understandings in a physics classroom. *International Journal of Science Education 19* , 209-228.

van Zee, Emily, & Minstrell, Jim. (1997b). Using questioning to guide student thinking. *Journal of the Learning Sciences, 6* (2), 227-269.

Wittrock, Merlin C. (Ed.). (1986). *Handbook of research on teaching* (Third edition). New York: Macmillan.

Yinger, Robert. (1977). *A study of teacher planning: Description and theory development using ethnographic and information processing methods.* Unpublished doctoral dissertation, Michigan State University, East Lansing.

Zimmerlin, Daniel, & Nelson, Mark. (1996). The detailed analysis of a beginning teacher carrying out a traditional lesson. Paper presented at the annual meeting of the American Educational Research Association, New York, April 8-12, 1996.

Appendix A

Students work in groups on:

(a) $\frac{m^4}{m^4}$ (b) $\frac{x^3 y^7}{x^2 y^6}$ (c) $\frac{x^5}{x^5}$

as T circulates checking student work. T then walks to front board and addresses class.

1	T:	OK. So on (a), what did we get on (a)?
2	Ss:	m to the fourth.
3	T:	m to the fourth.
4	T:	Anybody not get that?
5	S1:	Me.
6	T:	You got it. I saw it on your paper, [S1]. Don't you give me that.
7	T:	OK, (b)?
8	Ss:	xy . xy .
9	T:	OK. Why did you do that? Why did you get xy ?
10	Ss:	[inaudible.]
11	T:	[raising hand] Hold it. One person. One person.
12		[Class quiets and S2 raises his hand.]
13	T:	[S2], why did you get xy ?
14		[S2 starts to answer, quietly, and class noise overrides.]
15	T:	Shhh.
16	S2:	[inaudible.]
17	T:	OK. What did you subtract?
18	S2:	Uh, 3 minus 2.
19	T:	OK. You looked at the x 's [pointing to x -terms in numerator and denominator] and [pointing to exponents] you subtracted 3 - 2. That gave you x to the first [writes x on the board].
20	S2:	Yup.
21	T:	And then [points to y terms] you looked at the y 's and said [points to the exponents] 7 minus 6, gives you y to the first [writes y on board].
22	T:	All right. [Looks around at students. S2 gives thumbs up.] OK.
23	S3:	[inaudible.]
24	S2:	[talking over S3] x to the zero.
25	S3:	. . . mine don't look like that.

26 T: OK, (c). What did you get here?
27 Ss: [overlapping] x to the zero. Nothing. Just zero. One.
28 T: [pointing to various students] One. x to the zero.
29 OK. [Writes equal sign (=) and a fraction bar after problem (c).] Let's work this out. Let's work this one out. [pointing to numerator] This says how many x 's do we have?
30 Ss: Five.
31 [T writes *xxxxx* on board in numerator.]

32 T: And how many x 's are on the bottom?
33 Ss: Five.
34 T: [writes *xxxxx* on board in denominator] OK.
35 T: So what do I do?
36 Ss: Cancel.
37 T: [cancels]
38 T: So what am I left with?
39 Ss: [overlapping] x . Zero. x . [inaudible.]
40 T: There's nothing here? Is there zero here?
41 Ss: Yes. No. Zero over zero. Nada
42 T: [holding up hand in gesture to wait] Just a second. Just a second.
43 T: [moves to blank section of board and writes] What if we had 5 over 5? What's that equal?
44 Ss: One.
45 T: But we, I just cancelled them.
46 S4: Nothing.
47 S5: But there's a 1 there.
48 T: [pointing to 5/5] But there's a 1 there. [pointing to *xxxxx* /*xxxxx*] Is there a 1 there?
49 Ss: [overlapping] Yes. No.
50 T: Yeah there is.
51 S6: There ain't no 1 there.
52 [T visibly slumps and steps back. Waits.]
53 Ss: [overlapping] Yes there is. [inaudible.] x to the zero power is zero. [inaudible.]
54 S7: x .
55 T: It's just x .
56 S7: Yeah.
57 T: But there's no x 's left.
58 S7: You still got them right there.
59 S5: Wait [inaudible].
60 T: But they're cancelled out.
61 Ss: [inaudible.]
62 T: Shhh. Hold on. Hold on.
63 [Class quiets.]
64 T: OK. So I cancelled all these out [making cancelling motions over x 's]. Is it, is this zero over zero?
65 S8:
66 S9: No, it could be . . .
67 S5: It's 1.
68 Ss: [inaudible.]
69 T: What is it? What do you think?
70 S2: x .
71 T: But there's no x 's left.
72 S7: Yes there *is*.
73 S10: No it's *not*.

74 T: Shh. Hold on. Hold on. [S5]?
75 S5: x to the zero power over 1.
76 T: [with finality] OK. Very good. Five over 5 equals 1. So this [pointing to *xxxxx* /*xxxxx*] is going to equal 1 as well.
77 [S2 raises hand.]
78 T: Yes?
79 S2: x to the zero power over 1.
80 T: [to S2] Right. [to class] It equals 1. Or if you subtracted them what did you get? Someone said this earlier. x to the what power?

81 S5: Zero.
82 T: x to the zero. x to the zero equals what?
83 S5: One.
84 S11: One! I knew it!
85 [T writes Zero Exponent, $x^0 = 1$ on board and puts a box around it.]
86 S7: [reads] Zero exponent.
87 T: OK. Get this in your notes. Zero exponent. x to the zero - any number to the zero is 1. I'm going to put a [changes it to $a^0 = 1$]. It could be *any* number. a to the zero equals 1. Any number to the zero power equals 1. Zero power means exponents are the same over each other. OK. OK.

APPENDIX B
Minstrell "Best Value" discussion, Chunk [1.2.2]
Taped and transcribed by Emily van Zee; revised by Alan Schoenfeld

T: All right. Now. We're trying to get a best value and we might take all of the numbers or we might take some of the numbers and then it's, what the heck are we going to do with those numbers? OK?

So now we've got some numbers there, what are we going to do with those numbers? What's one thing that we might do with the numbers? [S5]?

SS: Average them

T: OK. [writes "average them" on board] We might average them. Now what do you mean by "average" here, [S5]?

S5: Add up all the numbers and then divide by whatever amount of numbers you added up.

T: All right. That is a definition for "average."

In fact, that's what we'll call an "operational definition." An operational definition is a definition where you, where you give a recipe for how to find what it is that you're talking about. And in this particular case, she's saying, "Add the number of - whether you're talking about some of the numbers or all of the numbers - add those up and divide by however many there are." And that's called the arithmetical average and to get that you add them up and divide by how many there are. OK?

[talking while writing this on the board] That's an average that you often use in lots of different contexts and it's an average that we'll use in here,

but look out! because there are lots of times when that's not the best average to use. On finding the best value, that's a pretty good way to get the arithmetic, or the arithmetic average is a pretty good way of getting a best value. OK?

Any other suggestions for what we might do? So we can average them -

[8s pause]

Any other suggestions there for what we might do to get a best value?

I'll put up the numbers that we had from the table measurement on Friday in first period, for the length, for the width and the length. As you look at that array of numbers, any other ideas there that come to mind as to how you might go about getting a best value from these numbers you're going to take there? [S7]?

S7: You've got a bunch of numbers that are the same number.

T: OK. Like what are you talking about there?

S7: 107

T: All right. 107 point zero, 107 point zero, 107 point zero, 107 point zero.

Is there any other number in the width column that shows up as much as 107 point zero?

S: No.

T: No. OK? So it's the number that shows up the most often is another way of picking one. That's called the, um, the mode. OK? [writes on board] The number that shows up most frequently. OK. That's another way of getting a best value out of a collection of numbers that you're willing to keep.

Does that make sense? Anybody confused here? yet? Haven't confused anybody yet? Then I've got to push a little harder. [4s pause]

All right?

Anybody think of another way of giving a best value? [S5]?

S8: This is a little complicated but I mean it might work. If you see that 107 shows up 4 times, you give it a coefficient of 4, and then 107.5 only shows up one time, you give it a coefficient of one, you add all those up and then you divide by the number of coefficients you have.

T: You lost me.

S?: [unintelligible] [overlapping student comments]

S8: One of those numbers. It's just that the more times it shows up, that makes like makes it a more, um, a more weight.

T: OK. Let me see if I can follow what you're saying then. You're saying one zero seven point zero shows up four times [writing on board] so let me put a multiplier in front of it, [sotto voce: that's what a coefficient is], of four, and then what, what am I going to do?

S8: Ah, you average that, well then you, just say there are, ah, five numbers, and another one is

T: Well let's go ahead and use this first column right here.

S8: OK. Then, ah, well, [unintelligible]

T: So everything else only comes up once

S8: Wait. one, yeah, looks like it. So everything else just gets one.

T: All right. So one and, ah, we've got, ah, one oh six point eight.

S8: Eight.

T: And one?

S9: Oh seven point five

S8: Oh seven point five

T: One oh seven point five. And one?

S?: Six point five

T: One oh six point five.

S?: One oh six.

T: One, one of six. OK. Now what do I do?

S8: You add all that.

T: OK.

S8: And you divide by, [muttering], eight.

T: One, two, three, four, and four makes eight?

OK?

[Instructor has written on the board:

$$\frac{4\,(107.0) + 1(106.8) + 1(107.5) + 1(106.5) + 1(106.0)}{8}$$

T: All right. What do you think of that method?

Ss: [overlapping student comments including "Forget it." "Too hard."]

T: Too hard?

Ss: [overlapping unintelligible student comments including "It's the same"]

T: All right. So actually it ends up being the same as the arithmetic average?

S8: No. Because 107 gets four times the value, so the 107 counts more.

T: Ah. OK. If you were to take the arithmetic average of these numbers, what would you do? What would be the operations that you would go through there?

[S5], you were the one who suggested arithmetic average

S5: You'd add all the numbers together and then divide it by 8.

T: Now what do you mean by "adding all the numbers"?

S5: You would add each separate number that everybody got; you wouldn't just add one 107, you'd add all the 107s.

T: OK. All right. So what [S5] is suggesting is for an arithmetic average is to add this number, then add this number, then add this number, even though it's a repeat of that one, then add this one, this one, and this one?

Ss: [overlapping unintelligible student comments]

T: Now would that come out the same as this if you did this?

Ss: [overlapping Yeah, yes, it would]

T: All right. So if you just took the arithmetic average by adding each one of these numbers, all eight numbers, and divide it by eight then, that would end up giving you the same number as this, so this is kind of maybe a quickie way of grabbing some of them, but outside of that, it it gives us the same answer?

S8: Yeah. It does. I didn't mean it too when I did it though.

T: OK. What about this other method that, ah, that was mentioned, of saying, let's just add up the numbers that are different? like 106.8 and 107.0, 107.5, 106.5 and 106.0, that's all our different numbers, right?

S9: Why 107.0?

T: Well, because that's a, that's a different

S10: It's different

T: I mean, there's at least one of those, at least one of these, at least these, etcetera, add those up and then take that and divide by 5. How do you like that?

S?: No.

S?: That doesn't show, doesn't represent it truthfully though, 'cause, I mean, there's a lot more 107s and that'd be, that'd change

T: OK. Wouldn't that give us the same answer as if we just took the arithmetic average?

Ss: [overlapping "no"s]

T: Can you give me an instance that's a real clear example, that would drive home to me as to why that would give me a different number than if I took all of these and divided by eight? We can do that with the numbers and see that it would come out different. [S11]?

S11: If everybody got 107 except for one person who got 99 and then if you took 107 and 99 and divided it by 2, it'd be a lot different.

T: Does that make sense?

Ss: [overlapping agreement]

T: So if people go back there and measure the table, 107,107, 107, 107,107, 107, 107, 107, 107, 107, 107 and somebody else gets 99, so we go over there and we say, Hmm, ah, half way between 107 and 99

Ss: [unintelligible comments]

T: Does that make sense?

Ss: [overlapping comments, "no"]

T: Now I want you to listen to yourselves because a lot of you are saying,

"that's a ridiculous situation; of course, it wouldn't be half" - what is half way between 99 and 107?

S?: 103

T: 103?

S?: Yeah.

T: OK. Of course it couldn't be 103, right? But you know what? There are going to be some contexts within here in which some of you are going to fall into that very trap right there, if you're not careful. OK? So watch out, so watch out for it. Is it clear that one oh, what'd I say, 103 would not be a good average for 99 and then all those 107s? Is that clear? OK. All right. Ah. OK. So this is really not a very good way to do it. Do we agree there? Somehow we need to weight, to weight in there the fact that 107 occurs so many times. So we've got this way of doing it, or if we

added them all up in there, that would include all those 107s. OK? [4s pause]

Anybody confused yet? [2s pause] No? [2s pause]

OK. [3s pause] Got to be honest. [4s pause]

All right. Anybody else see a different way of approaching?

S12: You could possibly take the number that appears most often, like you were saying before, if everyone got 107, and then a couple of people got 99, or like one person got 99 and one person got 120, you could pretty much assume that 107 would be nearest to the correct answer and [so that you could just select that]

T: OK. And that's the one that we called the mode there; it shows up the most frequently. OK. The mode.

There's another measure in here that, ah, that, ah, is sometimes used and nobody mentioned it but I'll, I'll go ahead and throw it in here then; it's what's called the "median" measure. Anybody know what the median is?

S?: Half way

S?: Half

T: Yeah. If you were to take, if you were to take all of the numbers - this is getting pretty messy there, let me clean that up a bit - if we were to take all of these numbers and rank them, [writing on board] the highest one is 107.5, then it's 107, 107, 107, 107, and then 106.8, 106.5, 106.0, do you see what I did there?

S?: Um hum.

T: I, what's called ranked them, from the biggest measure that we got to the smallest measure that we got for the width of that table, and then after I rank all of those, I go for the middle number, the middle number. Oh Beep, what do I do here? The middle number right in there.

S?: They're the same number so it doesn't matter.

T: Is that zero then? 'Cause it's right in there? [pointing between two numbers]

Ss: [overlapping comments: no]

T: Nah. 107's above it, 107's below it, right in there, I might even go half way between these two if they differed maybe, but the median number in this case would probably be a nice 107. OK? So the median is "the [writing on board] middle number when all are ranked." And ranked, you know, like from the top to the bottom, etc., so then you take the middle number when all the numbers are ranked there. You have to rearrange all the numbers and then take the middle one. And that's called the median. That make sense?

S?: Sure.

T: OK. All right. Now those are some of the, ah, some of the ways then that we might, we first of all might take all of the numbers to get a best

value, or we might take some of those numbers to get a best value, then what we might do with them is that we might average them, or we might go for the number that shows up most frequently, or we might go for the middle number. These are all different techniques for getting a best value.